UNIVERSAL COMMUNICATION VIA ROBUST COORDINATION

JACOB D. LESHNO AND MADHU SUDAN

Abstract. We ask whether two strangers are able to learn to meaningfully communicate through repeated interaction. We focus on a simple and essential part of meaningful communication: the two players must coordinate on common interpretation of messages. We formulate this problem as a repeated coordination game and ask whether the two players can guarantee that after a finite learning time they are able coordinate in all periods forward. We show that a "grain of coordination" can be leveraged to eventual coordination, but it is impossible to achieve coordination deterministically without some initial coordination. When players are initially symmetric eventual coordination can be guaranteed only with the use of randomization, implying that the common established encoding rule must be randomly chosen.

1. Introduction

A keystone in the theory of computation is the celebrated Universal Turing Machine — there exists a single computer capable of executing programs written in any language. The Universal Turing Machine allows anyone to receive and successfully run programs from any stranger. We ask whether we can achieve analogues universality in the context of communication. Under what conditions a single communication device will be able to meaningfully communicate with any stranger device?

A key challenge is that meaningful communication requires more than just the exchange of signals, it also requires coordination on a common language that gives an agreed interpretation of signals. Suppose two strangers start their

Leshno: Columbia Business School, Uris Hall, Columbia University, New York, NY 10027, USA. jleshno@columbia.edu. Sudan: School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA. madhu@cs.harvard.edu. We would like to thank Navin Kartik, John Nachbar, Larry Samuelson, Dov Samet and two anonymous referess for helpful comments and suggestions. Part of the work was done when the authors were in Microsoft Research New England.

interaction without a common language, can they attain a common language through repeated interaction?

Formally, we consider two players, Alice and Bob, who play a repeated coordination game with perfect monitoring, but in a changing environment and with limited prior knowledge about each other. We ask whether they can guarantee **eventual coordination**: They would like to guarantee that after some finite number of periods they establish an agreed rule that will allow them to coordinate in every round. In the context of Example 2 below, the players reach eventual coordination if they are able to establish a common language through repeated play. We use the term **universal strategy** to describe a strategy which guarantees eventual coordination.

We start by considering the asymmetric problem from Alice's perspective: Alice wishes to guarantee eventual coordination given minimal knowledge on the strategy played by Bob, described by a (possibly large) set O of possible repeated game strategies for Bob. We ask: Is there a universal strategy for Alice which will guarantee eventual coordination with any strategy from O? We show that under two restrictions on the set O the answer is positive. In particular the answer is positive if Alice can assume that Bob is attempting to coordinate and can only play Turing computable strategies.

We then proceed to ask if the existence of the universal strategy can allow Alice and Bob to reach eventual coordination in a symmetric setting. Without any prior coordination both Alice and Bob must have symmetric roles and symmetric knowledge. They both know that they both play strategies in O; can they guarantee eventual coordination while keeping this knowledge true? We show the answer is negative when the actions are not a priori distinguishable. In fact, any deterministic universal strategy for O does not belong to O. An immediate and important corollary is that while there exists a universal strategy that can guarantee coordination if the other player plays some deterministic and Turing computable strategy, any such deterministic universal strategy is not Turing computable.

We interpret the results to say that eventual coordination requires "a grain of coordination". When there is a slight asymmetry between the players it can be leveraged through repeated play to eventual coordination, as the asymmetry allows us to designate one player to be an "active learner" while the other is "passive" and eventually learned. But when the players are a priori symmetric in terms of their roles, then repeated play cannot guarantee eventual coordination. While the proof is extremely simple it highlights the advantage of the simple setup.

We follow to consider randomized strategies. We find that randomization allows us to reverse the impossibility results and attain coordination in symmetric settings as well. To give some intuition, if Alice and Bob could coordinate on one common coin flip, they could break the symmetry and agree who is an active learner and who is passive. If they randomize independently there is a positive probability that they will manage to replicate the common coin and successfully break the symmetry. By using strategies such that never lose coordination once it is reached, and are able to keep attempting while coordination was not reached, we give universal strategies that achieve eventual coordination from initial randomness. These strategies attain a "grain of coordination" through randomness, and leverage it to eventual coordination.

Put together, our results imply that universal communication is possible, but requires randomization. The randomized universal strategies we present allow two agents with minimal knowledge about each other, for example a limitation on available computational power, to establish a common language. The randomization is necessary and implies that the common language reached must be selected at random.

1.1. **Related and prior work.** This paper is a part of a large literature studying the economics of language, starting with Marschak (1965). This literature explores the connections between the structure of language and decision making (Rubinstein (2000)), structure of the firm and language (Dessein (2002), Cremer et al. (2007)) and experimentally investigates how subjects learn to communicate (Blume et al. (1998), Selten and Warglien (2007)). Our paper contributes to this literature by showing that language can be learned from very little structure. We show that initial asymmetry facilitates the learning of a common language, and without asymmetry randomization is necessary for learning.

Our work is inspired by the work of Goldreich et al. (2012) (see also Juba and Sudan (2008, 2011)) who studied meaningful communication in the context of solving complexity-theoretic oriented computational goals. The simplicity of our setting allows us to obtain clarifying insights on the necessary requirements for universality which were hard to isolate from other computational restrictions in the previous settings.

Schelling's seminal work (Schelling (1980)) explored how focal points provide basis for coordination, even in one shot games with plethora of possible options. While in many settings agents may have some prior coordination, there are many settings in which agents find it hard to coordinate. For example, Crawford et al. (2008) find that slight perturbations can annul the power of focal points. Our paper assumes no initial coordination, and explores how coordination (or focal points) is learned.

Our paper also follows the work of Crawford and Haller (1990) which studies optimal play in coordination games where the lack of common language restricts players to symmetric strategies. Our model differs in two ways. First, we allow the setting to change from period to period. While in Crawford and Haller (1990) a single round of coordination is enough to establish coordination in all future periods, in our model the players need to establish a rule or common language, and a single round of coordination may not suffice. For instance, in example 1 below coordination on a restaurant in Denver will not inform the players which restaurant they should choose in Salt Lake City. Farrell (1987) consider the battle of the sexes and shows that pre-game cheap talk does not resolve the coordination problem. Following works by Bhaskar (2000); Kuzmics et al. (2014); Sandroni (2000) study repeated games and show that the restriction to symmetric strategies limits the set of attainable payoffs. Blume (2000) shows that if agents have stronger knowledge of each other they can learn to coordinate more efficiently.

Kalai and Lehrer (1993) study convergence to equilibrium through repeated play. They study general games and show that if the players have a prior belief about each other's strategy that has a "grain of truth" (in that both players put positive probability on each other's strategy) then they converge to a path of play that will be approximately consistent with Nash equilibrium. While this work has many similar ingredients to our study, their convergence may be to a mixed equilibrium and does not guarantee coordination. Nachbar (1997) and Nachbar (2005) prove impossibility results showing that the Kalai and Lehrer (1993) convergence would fail under some natural requirements. We obtain similar negative results for deterministic strategies, but differ in that we get positive results when we allow randomization.

Our solution concept asks for strategies that will work well even under a "worst-case" assumption. Bergemann and Morris (2005) propose similar "robust" solution concepts and ask what can a mechanism designer do under minimal assumptions on type spaces, i.e. preferences of players and their belief on other players. Xandri (2012) uses a similar solution concept to study reputation building.

1.2. **Organization of this paper.** Section 2 provides examples illustrating the coordination problem. In Section 3 we introduce our model formally and present our formal definition of eventual coordination. In Section 4 we consider the asymmetric setting and present some positive results. In Section 5 we consider the symmetric setting and present our negative result for deterministic strategies. In Section 6 we consider the symmetric setting and present positive results using randomness. Concluding remarks are presented in Section 7.

2. Communication as coordination

In this section we present two examples to illustrate our ideas. The first is a variation of a classic repeated coordination game. The second is a repeated sender-receiver communication game. In both settings the agents can easily win if they pre-coordinate before the game. We ask whether two players with limited prior knowledge of each other can establish the same level of coordination through repeated interaction.

First, consider the following repeated coordination game:

Example 1. Alice and Bob are two business persons who often travel to the same conferences. Each conference is in a new city, so the conference brochure always lists two suggested restaurants. Alice and Bob simultaneously choose where to eat. Both restaurants are equally good, and their preferences are to eat at the same restaurant.

Each round of the game in example 1 is a coordination game. Each city entails a new set of options, but the description of the two options is common knowledge. If Alice and Bob agree to follow a rule to determine where they meet, such as "choose the restaurant which is closest to the hotel" or "choose the restaurant whose English name is alphabetically first", they would be able to coordinate and eat together at every conference. But suppose that Alice and Bob have limited prior knowledge of each other, can their repeated interaction enable them to establish such a rule so that they will coordinate and eat together from some conference onwards?

In our second example Alice and Bob play a communication game, which is given to illustrate how coordination is essential for meaningful communication:

Example 2. Alice and Bob are randomly selected contestants on a game show. Each round a prize is hidden behind one of two curtains, and both players win the prize if Bob selects the correct curtain. Alice sees which curtain holds the prize, and can pass a single message to Bob as follows. Alice is given two pictures, and selects the order in which they will be presented to Bob. After observing the order, Bob picks a curtain. The round ends with both players learning whether they won, and two new pictures are drawn for the following round.

In example 2 each round is a sender-receiver game, with two possible messages. The challenge is that although Alice can always send one bit of information (one of two possible messages), she will manage to transmit the intended meaning to Bob and win the prize only if they coordinate on the same interpretation of messages. Let the pictures of the round be X and Y, and denote the two possible messages by $\langle X, Y\rangle$ and $\langle Y, X\rangle$. There are two informative encodings: under encoding-I $\langle X, Y\rangle$ means "prize is behind left curtain" and $\langle Y, X\rangle$ means "prize is behind right curtain"; under encoding-II $\langle X, Y\rangle$ means "prize is behind right curtain" and $\langle Y, X\rangle$ means "prize is behind left curtain". Assuming that Alice and Bob are both trying to communicate, they will each choose either encoding-I or encoding-II, and win this round if and only if they both choose the same encoding. Winning every round requires them to be able to agree on an informative encoding for any pair of pictures, which is possible if they establish a "common language".

This simplified setting illustrates how coordination is essential for communication. To win a round in example 2 it is not enough for the two players to communicate a bit of information, they need to **meaningfully communicate**, which in turn requires them to coordinate and agree on the meaning of messages. In typical communication setting the agents are indifferent between the different encodings, as long as both sides are in agreement. Thus, solving a coordination game is essential for meaningful communication, in the sense that if Alice and Bob wish to win in example 2, they need to be able to win the coordination game of example 1.

Note that in example 2 the possible messages change from round to round, but the message space in each round is common knowledge. If Alice and Bob come into the game with prior agreement on a "common language", or a rule for interpenetrating messages, they can win every round. For example, one such rule for interpreting messages can be "ordering pictures according to English lexicographic order means the prize is behind the left curtain, and the reverse order means the prize is behind the right curtain".[1] But if Alice and Bob have limited prior knowledge of each other, can they learn how to meaningfully communicate through their repeated interaction? Can they establish such a rule that will allow them to win every round forward?

In the following we consider the coordination problem that is inherent to both example 1 and example 2. In both examples the incentives of the two players are aligned, and once they reach agreement on a "rule" they will win every following round. But suppose the Alice and Bob have limited prior knowledge of each, and entering the game of example 2 Alice uses the rule "lexicographic in English means right" but Bob uses the rule "lexicographic in Urdu means right". While they may be able to win some rounds by chance,[2] one of them will have to switch rules if they want to win every round from some point. Alice should play strategically to try to match the same rule as Bob, taking into account that Bob may be switching to try and match Alice.

We ask whether Alice and Bob can guarantee **eventual coordination**: Alice and Bob repeatedly play a coordination game, with limited prior knowledge about each other, and they would like to guarantee that after some finite number of rounds they establish coordination in every round forward. In the context of example 1, the players reach eventual coordination if they are able to

[1]Another example is the following rule that avoids the ambiguity of the English language: "convert the pictures to binary vectors according to the BMP file format; a lexicographic order of the vectors means that ...".

[2]Note that unlike in Crawford and Haller (1990), a single round of coordination is not sufficient to establish coordination in all following periods.

agree a rule for choosing restaurants. In the context of example 2, the players reach eventual coordination if they are able to establish a "common language" through repeated play. We use the term **universal strategy** to describe a strategy which guarantees eventual coordination with the other player. We ask whether a universal strategy is possible, and if it requires initial asymmetry or randomization.

3. Model and Definitions

We use $i \in \{1,2\}$ to denote a player and $t \geq 0$ to denote a period. In each period t player i observes a description $\theta^t \in \Theta^t$ and chooses between two of two possible actions $\hat{a}_i^t \in \hat{A}^t$. For ease of description, we arbitrarily choose a one-to-one and onto mapping $\mathcal{L}_0 : \Theta^t \times \hat{A}^t \to A = \{0,1\}$, relabeling the actions as $a_i^t \in A = \{0,1\}$. We emphasize that these labels are arbitrary, and we do not assume that the labels of the actions are common knowledge. The action profile of period t is $a^t = (a_1^t, a_2^t) \in A^2$. The players coordinate and receive a payoff of $+1$ in period t if $a^t \in \mathcal{C} = \{(0,0)\,(1,1)\}$, and receive a payoff of 0 otherwise.

Each player observes the history of actions played,[3] as well as the descriptions of previous periods. Denote the history at the beginning of period t by $\hat{h}^t = \{(\theta^\tau, a_1^\tau, a_2^\tau)\}_{\tau<t} \cup \{(\theta^t)\}$, and let $\hat{H}$ be the set of all possible histories. Throughout our analysis we will hold $\{\theta^t\}_{t\geq 0}$ fixed, and denote by $H = (A \times A)^*$ to be the set of histories under the labeling $\mathcal{L}_0$. We denote player i's strategy by $S_i : H \to \Delta(A)$. With slight abuse of notation, we refer to deterministic strategy S_i as a mapping $S_i : H \to A$.

Remark. A strategy encompasses both the player's strategic decisions, as well as the player's subjective labeling of the actions. As an illustration, consider Alice in example 1, and assume she wants to follow a strategy which we loosely describe as "switch restaurant if not coordinated last period". If Alice is using a labeling rule $\mathcal{L}_{Alice}$, she may define "switch restaurant" in period t as taking the a_i^t such that $\mathcal{L}_{Alice}(\theta^t, a_i^t) = 1 - \mathcal{L}_{Alice}(\theta^{t-1}, a_i^{t-1})$. If Alice's labeling is different from our encoding of the game, that is $\mathcal{L}_{Alice} \neq \mathcal{L}_0$, we can write

[3] If players observe their payoff they can infer the other player's action.

her strategy S_i under our labeling $\mathcal{L}_0$; The strategy S_i will depend on Alice's labeling.[4]

To describe our results it would be useful to describe strategies as state machines, as defined by Mailath and Samuelson (2006); under this notation, a strategy for player i is described by a pair of functions

$$\begin{aligned} S_i : &\quad \Sigma_i \to \Delta(A_i) \\ \Phi : &\quad \Sigma_i \times A^2 \to \Sigma_i \end{aligned}$$

and an initial state $\sigma_i^0 \in \Sigma_i$. Σ_i is the (not necessarily finite) state space of S_i. If the state is $\sigma_i^t \in \Sigma_i$ at the beginning of period t, then $a_i^t = S_i(\sigma_i^t)$ at time t. After observing the joint action profile a^t the next state is set to $\sigma_i^{t+1} = \Phi(\sigma_i^t, a^t)$.[5] Denote by $S_i(\sigma_i)$ the strategy S_i starting from the state σ_i. A strategy profile $S = (S_1, S_2)$ gives a probability distribution $\mathbb{P}_S$ over H, describing the probability of future paths of play.

The player's goal is to guarantee eventual coordination:

Definition 1. Two strategies S_1, S_2 **eventually coordinate** if the induced path of play includes only a finite number of non-coordination periods with probability 1. That is,

$$\lim_{T\to\infty} \mathbb{P}_S(\{h \in H \mid h^t \in \mathcal{C} \ \forall t > T\}) = 1.$$

In examples 1 and 2, two players that come into play with an agreed upon rule can win every period. If the players come into the game with limited prior knowledge of each other they should be able to win in the first period, but

[4]For example, Alice may use the labeling rule $\mathcal{L}_{Alice} = \mathcal{L}_{close}$ which is "the closer restaurant is labeled 0", while we labled the strategies by the labeling rule $\mathcal{L}_0 = \mathcal{L}_{alphabetic}$ which is "the restaurant which is first alphabetically in English is labeled 0".
If Alice's deterministic strategy is $\tilde{S}_i : \tilde{H} \to A$ where $\tilde{H}$ is the set of histories under her labeling $\mathcal{L}_{Alice}$, her strategy can be translated to $S_i : H \to A$ under our labeling $\mathcal{L}_0$ by the following translation

$$S_i(h^t) \triangleq \mathcal{L}_0 \circ \mathcal{L}_{Alice}^{-1} \left(\theta_t, \tilde{S}_i \left(\mathcal{L}_{Alice} \circ \mathcal{L}_0^{-1} \left(\tilde{h}^t\right)\right)\right)$$

where $\mathcal{L}_{Alice} \circ \mathcal{L}_0^{-1} \left(\tilde{h}^t\right) \triangleq \left\{\left(\mathcal{L}_{Alice} \circ \mathcal{L}_0^{-1} (\theta^\tau, a_1^\tau), \mathcal{L}_{Alice} \circ \mathcal{L}_0^{-1} (\theta^\tau, a_2^\tau)\right)\right\}_{\tau<t}$, and we slightly abuse notation by writing $\mathcal{L}_{Alice} \circ \mathcal{L}_0^{-1} (\theta^t, a^t)$ instead of $\mathcal{L}_{Alice}\left(\theta^t, \mathcal{L}_0^{-1} (\theta^t, a^t)\right)$.
[5]This extends the previous definition, as we can set $\Sigma_i = H$ and $\sigma^t = h^{t-1}$. Furthermore, in the absence of computational restrictions on the strategies, the two definitions can be shown to be equivalent.

we ask that with probability 1 after some finite number they reach agreement that allows them to win every period forward.

We model the player's initial knowledge as a set O of possible opponent's strategies, if $S_j \notin O$ each player is certain that the opponent is not playing S_j. The game form and the descriptions $\{\theta^\tau\}_{\tau \le t}$ are common knowledge. The labeling $\mathcal{L}_0$ is not assumed to be common knowledge, and we can capture uncertainty about the encoding used by player j by including multiple strategies in O.[6]

Definition 2. Player i with knowledge O can **guarantee coordination** if there exists a strategy U such that for every $S_j \in O$ the strategies U, S_j eventually coordinate. We refer to such U as a **universal strategy for** O**.**

Note that if U is universal for O and $O' \subset O$ then U is universal for O' as well. We will try to let O be as large as possible, and exclude strategies by making minimal assumptions which are necessary for universality.

Remark. In the communication game of example 2 players have four actions in each period: two informative encodings, and two non-informative (Alice chooses one of the two possible messages and sends it regardless of where the prize is, Bob ignores the message and always opens one of the two curtains). Establishing a common language in the game is at least as hard as eventually coordinating in our coordination game, since meaningful communication requires both players to use informative encodings, and after excluding strategies that use non-informative encodings we are left with a two action coordination game. Also note that in terms of our results, a player that uses a non-informative encoding is not materially different from the player using a random action that gives equal probability to each informative encoding.

4. Asymmetric Universal Coordination: positive results

We start by considering the coordination problem from player 1's perspective. Player 1 knows that player 2 will play some strategy in O and wishes to guarantee eventual coordination. Before stating sufficient conditions for universality, we describe some of the obstacles. We start with a simple example

[6]The See the remark above.

that shows that there exists strategies whose containment in O prevents the guarantee of eventual coordination.

Example 3. Let S_{mix} be the strategy that plays 0 with probability $\frac{1}{2}$ and 1 with probability $\frac{1}{2}$, independent of the history.

It is clear that no strategy S_1 for player 1 will eventually coordinate with S_{mix}. If O includes any strategy that does not allow eventual coordination with some strategy then there can be no universal strategy for O. Therefore we need to require that O includes only strategies which some strategy can coordinate with. We stress that the quantifiers are "switched": Our requirement here is that for every strategy $S_j \in O$ there is some strategy $C = C(S_j)$ that achieves eventual coordination with S_j. In contrast the condition for universality is that there should exist some strategy U such that U achieves eventual coordination with S_j for every $S_j \in O$.

Our next example shows that universality may be prevented not because of any single strategy in O, but rather because strategies conflict with each other. Consider the following example:

Example 4. Let $S^p_{\text{password1try}}$ be a strategy defined by a ten bit sequence $p \in \{0,1\}^{10}$; If the other player plays the sequence p in the first ten periods then $S^p_{\text{password1try}}$ plays 0 forever,[7] else it mixes 50-50 forever.

$S^p_{\text{password1try}}$ eventually coordinates with the strategy that plays p in the first ten rounds and plays 0 forever after. However, if O includes two strategies $S^p_{\text{password1try}}$ and $S^q_{\text{password1try}}$ with different passwords $p \neq q$, then there is no strategy which can coordinate with both. To rule out such incompatibilities that it is suffice to require that each strategy $S_j \in O$ can reach coordination after any history. Note that this is a condition on each strategy individually.

Definition 3. We say that a strategy $S_j \in O$ is **coordinateable**, if, for every possible state σ_j of S_j there exists a strategy $C = C(S_j, \sigma_j)$ such that C eventually coordinates with the strategy S_j started from the state σ_j. We say that O is **coordinateable** if every $S_j \in O$ is coordinateable.

[7] Note that "playing 0 forever" simply means playing $\hat{a}_i^t$ such that $\mathcal{L}_0(\theta^t, \hat{a}_i^t) = 0$, which is some fixed sequence of actions.

This brings us to our first requirement:

Axiom 1. *The set O is coordinateable.*

Note that every deterministic strategy is coordinateable. Unfortunately, even assuming that O is coordinateable turns out to be insufficient for universality, as demonstrated by the following example:

Example 5. A password strategy S_{password} is defined by an infinite sequence $\{p_k\}_{k=1}^{\infty}$ where each $p_k \in \{0,1\}^{10}$. S_{password} plays p_1 for the first ten periods. If the other player also plays p_1 in the first ten periods, then S_{password} plays 0 forever. Else, it moves to p_2 for the next ten periods, etc.

$S_{password}$ is coordinateable, a strategy that knows the relevant p_k will quickly coordinate with $S_{password}$. But if O includes all possible $S_{password}$ then there is no deterministic universal strategy for O. To see that, consider any deterministic strategy S_i and select a sequence of passwords $\{p'_k\}$ such that S_i misses each and every one of the passwords. A randomized strategy S_i which guesses passwords uniformly at random will reach coordination with probability 1 if all passwords are of constant length, but if we allow for password strategies such that $p_k \in \{0,1\}^{10k}$ even such a randomized strategy will not be universal for O.

Therefore we add our second requirement:

Axiom 2. *The set O is countable.*

The axiom above is even less immediate from the example preceding it, however, it is a natural restriction if we consider the complexity of computing a strategy. Suppose that player i first labels the actions using $\mathcal{L}_i$, and plays according to $\tilde{S}_i : \tilde{H}_{\mathcal{L}_i} \to \Delta(A)$ where $\tilde{H}_{\mathcal{L}_i}$ denotes the histories under the labeling $\mathcal{L}_i$. Each such strategy can be translated to into a equivalent description $S_i : H \to \Delta(A)$ under our labeling $\mathcal{L}_0$.[8] Suppose that player i can only use Turing-computable functions,[9] that is, both the labeling function $\mathcal{L}_i$ and the function $\tilde{S}_i$ are Turing-computable. Then the set O of all possible $S_i : H \to \Delta(A)$ strategies (described under labeling $\mathcal{L}_0$) is countable, because

[8]See footnote 4.

[9]A strategy is computable if the function S describing the strategy can be generated by a Turing machine (see Sipser (2006)).

each strategy is given by a combination of two Turing computable function. Thus, Axiom 2 will be satisfied if players are minimally restricted in their computational power. Note that when we translate between the different labeling we use the descriptions θ^t, that we assumed are common knowledge.

The following theorem shows that universality is possible under the two assumptions. In the theorem we restrict attention to deterministic strategies, showing that even without randomization player i can have a deterministic universal strategy. The proof can be extended to show that a deterministic universal U exists under the same condition even when O contains randomized strategies.

Theorem 1. *If O is a countable set of deterministic strategies, then there exists a deterministic strategy U that is universal for O.*

Remark. The proof below can be derived from the work of Goldreich et al. (2012), and included here for completeness.

Proof. We build U as follows. Let $\{S_j \mid j \in \mathbb{N}\}$ be an enumeration of the strategies in O. The universal strategy U proceeds in stages starting in stage 0. Let t_j denote the period at the beginning of stage j and let h_j denote the history of actions at the beginning of stage j. If h_j is inconsistent with the actions of S_j, then skip to stage $j+1$, else let ρ_j denote the state of strategy S_j after history h_j. Let $C_j = C(S_j(\rho_j))$ be an eventual coordination strategy for $S_j(\rho_j)$ (such a strategy exists since S_j is coordinateable) and let n_j be an upper bound on the number of non-coordination periods before C_j and $S_j(\rho_j)$ achieve coordination. U plays according to C_j till there are $n_j + 1$ non-coordination periods, and if this event happens, it moves to phase $j + 1$.

To verify that U is universal, suppose it is playing against S_j. We first claim that U never moves past stage j. This is obvious since if it reaches stage j and plays according to $C_j(S_j(\rho_j))$ then it will reach coordination with fewer than n_j non-coordination periods. Thus it follows that U stops in some stage $k \leq j$. We next claim that whichever stage it stops in implies coordination. Again this is straightforward since if U encounters more than n_k+1 non-coordination periods during stage k, it would to stage $k+1$. Thus U arrives to stage k in a finite number of periods and then only has a finite number of non-coordination periods before reaching perpetual coordination with S_j. □

The following corollary is immediate, using the fact that the set of computable strategies is countable. It gives a positive answer to our question, allowing player 1 to guarantee coordination while only assuming that the description in each period is common knowledge, and that player 2 has limited computational power and does not randomize.

Corollary 1. *There exists a universal strategy that guarantees coordination with every deterministic (Turing) computable strategy.*

In terms of example 1, Alice may not be sure whether Bob labels restaurants by whether they are close/far or alphabetically first/second. Alice may also be unsure whether Bob "switches" (under his labeling) after miscoordination or not. But Alice may know that Bob uses a deterministic computation to reach his decision, where the inputs are the descriptions and history of coordinations and the output is the chosen restaurant. If Alice can assumes a bound on Bob's computational ability she will be eventually be able to learn Bob's strategy and coordinate with him.

We warn the reader that the universal strategy itself may not be a computable strategy. In the proof above, C_j, n_j etc. need not be computable, which explains why our constructed strategy needs not be computable. We show a more serious obstacle in the next section: See the corollary to Theorem 2.

5. Symmetric Deterministic Strategies: Impossibility Results

The previous section concluded with the ability to construct universal strategies for a fairly wide class of knowledge sets. However, in the construction of that universal strategy we assumed asymmetry in the roles of the players: one player was asked to play any strategy within the knowledge set O and the other player was asked to play a universal strategy for O. Designating which player should assume which role is by itself a coordination problem. If the players coordinated on asymmetric roles, Theorem 1 allows them to leverage this to achieve eventual coordination.

In this section we ask whether the players can achieve eventual coordination without any prior coordination. We require that the players will have symmetric roles,[10] and hold the same knowledge set O on the other player's strategy. We therefore seek universal strategies for a set O that are themselves in O, which would allow the two players to take on symmetric roles in the attempt to reach eventual coordination.

It is possible to find knowledge sets O such that there exist a universal strategy that is itself in O, as illustrated by the following trivial example.

Example 6. Let $O = \{S_{\text{const0}}\}$ be the knowledge set which contains a single strategy S_{const0} which plays the action 0 after every history. Then $S_{\text{const0}} \in O$ is universal for O.

The set O hardly satisfies our initial goal to enable eventual coordination under minimal knowledge. Coordination in example 6 is not learned through repeated play; rather, play is coordinated from the first period. This requires that the action 0 is "special" , which in turns requires the players to have prior coodination that allows them to distinguish the two actions. Following this, we study knowledge sets where there is ambiguity on the labeling used by the other player.

We formalize this as follows. For $a \in \{0, 1\}$ we denote by $\overline{a} \triangleq 1 - a$ the label-switched action. For an action profile $a^t = (a^t_1, a^t_2)$, denote $\overline{a^t} = (\overline{a^t_1}, \overline{a^t_2})$, and for a history $h = (a^0, a^1, \ldots, a^t)$ denote $\overline{h} = (\overline{a^0}, \overline{a^1}, \ldots, \overline{a^t})$. For strategy S we define its **label-switched** strategy, denoted $\overline{S}$, to be the strategy that acts as S under label-switching: i.e., $\overline{S}(h) = \overline{S(\overline{h})}$. In other words, $\overline{S}$ acts the same as S, but acting with the labels of the actions switched, both on the history it observes and on the actions it outputs.[11] Note that the first action of $\overline{S}$ is the label-switched action of the first action of S.

Definition 4. A set of strategies O is **label neutral** if for every $S \in O$ we have $\overline{S} \in O$.

[10]This condition is akin to the positional-symmetry of Crawford and Haller (1990), and condition S of Nachbar (2005).

[11]Suppose we describe player i's strategy by a labeling function $\mathcal{L}_i$ and strategies $\tilde{S}_i : \tilde{H}_{\mathcal{L}_i} \to \Delta(A)$. The label-switched strategy is equivantly given by applying $\tilde{S}_i$ on histories labeled by $\bar{\mathcal{L}}_i$, given by $\bar{\mathcal{L}}_i(\cdot) = 1 - \mathcal{L}_i(\cdot)$.

The assumption that O is label neutral means that action labels are not meaningful on their own, and therefore the knowledge set O should remain the same if we switch the labels of the two actions. This condition is akin to action-symmetry and attainability of Crawford and Haller (1990) and condition C of Nachbar (2005).

We ask whether we can find a label neutral knowledge set O that will allow the two symmetric players to guarantee coordination using deterministic strategies. The following theorem shows that this is impossible:

Theorem 2. *Let O be a label neutral set of deterministic strategies. Then if U is universal for O then $U \notin O$.*

Proof. Note that when a deterministic strategy S plays against $\overline{S}$, both starting with the empty history, they fail to coordinate in every period. If $U \in O$ then we have that $\overline{U} \in O$ and they are both deterministic. Since U can not coordinate with $\overline{U}$ it follows that U cannot be universal. □

While the proof above is very simple, its ramifications are significant. Without some initial coordination on labels or asymmetry of roles, the players cannot guarantee eventual coordination. To illustrate this, consider again example 2. Suppose player 1 is passive - he is using a fixed encoding rule, say Japanese-alphabetical, and waits for the other player to adapt. If player 2 knows that player 1 is using a encoding rule, but does not know which, he can guarantee eventual coordination by using an "active" universal strategy: he attempts different encoding rules and eventually they coordinate on the same language. Theorem 1 shows that this "role asymmetry" allows the player to reach eventual coordination under a weaker condition on player 1's strategy, and Theorem 2 shows that the asymmetry is necessary. When players are symmetric they might both passively stay fixed and fail to coordinate, or both actively try to learn each other, and might keep switching and fail to coordinate. The initial role asymmetry serves as a "grain of coordination" which can be leveraged to eventual coordination, but without it eventual coordination cannot be guaranteed.

The following corollary helps us gain some intuition for the contrast between the negative result of Theorem 2 and the positive results of the previous section. Intuitively, to preclude that both players are cycling in attempt to learner each other, the universal strategy must be more powerful than any of the strategies it attempts to learn. Thus to guarantee coordination we must have asymmetry of computational power. The corollary shows that if the knowledge set is a class of strategies of bounded computational complexity, then the universal strategy requires more computational power than any of the strategies learned.[12]

Corollary 2. *A deterministic universal strategy that guarantees coordination with every deterministic Turing computable strategy cannot be Turing computable itself.*

This corollary can be extended to many other deterministic computational classes, including deterministic polynomial time, since all natural resource bounded classes are label neutral.

6. Symmetric Coordination: randomized possibility results

In Section 4 we saw that if the players can coordinate on asymmetric roles, which may be viewed as a "grain of coordination", it can be leveraged to attain eventual coordination. In this section we show that randomness can provide us with such a "grain of coordination". Our main result of this section is that under relatively mild assumptions on O, there exists a randomized strategy U that is universal for $O' = O \cup \{U, \bar{U}\}$. In particular, if O is label neutral then so is O'; so the existence of a universal strategy for O' in O' contrasts sharply with the deterministic impossibility result (Theorem 2).

Since we allow U to be randomized and ask for a symmetric setting, we must also allow for O to include random strategies, which U will have to coordinate with.

Theorem 3. *Let O be a countable set of (possibly randomized) coordinateable strategies. Then there exists a countable, label neutral set $O' \supseteq O$ of coordinateable strategies such that there exists a strategy $U \in O'$ that is universal for O'.*

[12]Knoblauch (1994); Nachbar and Zame (1996) show that there are repeated game strategies for which the best response strategy is not computable. Our setting differs in that the asymmetry in the player's computational is necessary and beneficial to both players. See also Gilboa and Samet (1989) and Kalai (1990).

This result is in contrast to Nachbar (1997) and Nachbar (2005) that obtain a negative result even under a class of randomized strategies. Our models differs in that we ask for eventual coordination rather than an approximate best response, and that we allow arbitrary randomized strategies.

We prove the result by showing that randomization can create the asymmetry which enabled our positive results in Section 4. To gain some intuition, suppose that the players could flip a common coin to determine roles: one player will play a universal strategy for O and the other player will play any strategy from O. Such a coin flip would allow the players to reach eventual coordination, as in Section 4. When the players do not have a common coin they can still privately randomize, each player separately flipping a coin to determine his role. With probability 1/2 the players select different roles and reach eventual coordination. With probability 1/2 the players fail to coordinate, but can try to flip coins again. Eventual coordination is guaranteed if coordination is an absorbing state and the players keep flipping coins while they are not coordinated.

Proof. Without loss of generality assume O is label neutral, else we can take $\tilde{O} = O \cup \bar{O}$.

We start by describing the strategy U that we later prove to be universal for $O' = O \cup \{U, \bar{U}\}$. Let $\{S_j \,|\, j \in \mathbb{N}\}$ be an enumeration of strategies in O in which every element of O appears infinitely often. Our universal strategy U runs in stages. At the beginning of stage j, U decides the actions of this stage probabilistically. With probability 1/2 it guesses that the other player is playing S_j and tries to coordinate with it as follows. Let h_j be the history at the beginning of stage j. If the probability of observing h_j when the opponent is playing S_j is zero, then we terminate the stage. Else let ρ_j be the state of S_j after history h_j. Let $C = C\left(S_j\left(\rho\right)\right)$ be the strategy that coordinates with $S_j(\rho_j)$ and let n_j be the minimal integer such that the probability that C achieves coordination with $S_j(\rho_j)$ before n_j miscoordinations is at least 1/2. U plays strategy C till it sees $n_j + 1$ miscoordinations and then moves to stage $j + 1$. With probability 1/4 U plays a constant 0s until there are $2N_j$ miscoordinations and with probability 1/4 U plays a constant 1s until there are $2N_j$ miscoordinations, where $N_j = \max_{j' \leq j, \rho_{j'}}[n_j(\rho_{j'})]$ is equal to maximal

n_j for any strategy $\{S_{j'} \mid j' \leq j\}$ and any state $\rho_{j'}$ that could be reached under the current history. If $2N_j$ miscoordinations happen, it moves to stage $j+1$.

To prove universality of U for O' fix a strategy $S_j \in O'$. We consider two cases: $S_j \in O$ and $S_j \in \{U, \bar{U}\}$. In the former case, we have that unless U gets stuck in some stage, it tries to coordinate with S_j infinitely often and for each attempt the probability of reaching coordination is at least has a $1/2$. On the other hand, the only reason U may get stuck in some stage is that it stops miscoordinating, so in either case coordination is reached with probability one in a finite number of periods.

Now we turn to the case where $S_j \in \{U, \bar{U}\}$. If U plays against $S_j = U$ or $S_j = \bar{U}$ then either coordination is reached, or it happens infinitely often that one of them initiates stage k' while the other strategy S_j is in stage $k < k'$. Wlog assume that U is at stage k'. With probability $1/4$ U plays constant 0 until $2N_j$ miscoordinations. Either coordination is reached within $2N_j$ miscoordinations, or S_j must start a new stage before U finishes its stage. If S_j starts a new stage there is a probability of $1/4$ that it plays constant 0 and coordination is reached. Thus whenever a new stage begins, the probability of reaching coordination is at least $1/16$, leading to eventual coordination with probability one in finite number of periods. □

The universal strategy U constructed in the proof alternates between trying to actively learn the other player, and passively playing a fixed sequence waiting for the other player to learn. If the set O is rich enough, when our strategy U is passive it is acting like a strategy in O, and should be learned by any strategy that is universal for O.

We remark that in the construction above several further restrictions are needed to make the strategy above computable. In particular the strategies in O should be computable, furthermore the set of coordinating strategies $C\left(S_j\left(\tau_j\right)\right)$ should be computable, and finally the number of miscoordinations n_j and N_j should be computable. Thus getting a computable universal strategy is non-trivial. In the following section we overcome all these restrictions by asking for a stronger notion of coordinatability, which allows us to relax our monitoring requirements.

6.1. **Untraceable states and uniform coordination.** Our model allows players to calculate the opponent's current state given an hypothesized strategy and history of play. The universal strategies we constructed so far required this ability, as the approriate response to a given strategy may depend on its state. We follow to strengthen our definition of coordinateability, and ask that we can coordinate with a strategy without knowing its current state.

Definition 5. S is **uniformly coordinateable** if there exists $C = C(S)$ and bounded function $k : (0, 1] \to \mathbb{N}$ such that for all states σ and for all $\epsilon > 0$, we have that the probability that C coordinates with $S(\sigma)$ with at most $k(\epsilon)$ miscoordinations is at least $1 - \epsilon$. We refer to C as the coordinating strategy for S.

Theorem 4. *Let O be a countable set of (possibly randomized) uniformly coordinateable strategies. Then there exists a countable, label neutral set $O' \supseteq O$ of uniformly coordinateable strategies such that there exists a strategy $U \in O'$ that is universal for O'.*

Notice that while Theorem 4 is weaker in that it requires the universal strategy to eventually coordinate only with uniformly coordinateable strategies, but it is stronger in that it requires U itself to be universally coordinateable.

Remark. If the set of strategies $\{C(S) \mid S \in O\}$ can be enumerated efficiently and computed efficiently, then the universal strategy can also be computable efficiently.

Proof. The universal strategy is similar to that of the previous proof, with some changes to ensure that the universal strategy itself is uniformly coordinateable, while exploiting that fact that the strategies in O are uniformly coordinateable. Again we assume O is label neutral, and take $O' = O \cup \{U, \bar{U}\}$.

We start by describing the universal strategy U: Let $c_0 \geq 1$ be some fixed constant. Let $\{S_j \mid j \in \mathbb{N}\}$ be an enumeration of strategies in O in which every element of O appears infinitely often. Let k_j denote the number of occurrences of S_j in $\{S_\ell | \ell \leq j\}$.

Again U runs in stages: At the beginning of stage j, U decides the actions of this stage probabilistically. With probability $1/4$, U plays 0s till there are c_0 miscoordinations, and if this event happens, it moves to stage $j + 1$.

With probability $1/4$ it plays 1s till there are c_0 miscoordinations. Finally, with probability $1/2$ it guesses that the other player is playing S_j and tries to coordinate with it as follows: Let $C = C(S_j)$ be the uniform coordinating strategy for S_j. U plays according to C until c_0 miscoordination steps occur, and then U tosses a coin. With probability $1/2$ U aborts this stage and moves to stage $j+1$, and with probability $1/2$ U continues the stage, and will toss the coin again if c_0 miscoordination steps occur.

To prove universality of U for O' fix a strategy $S \in O'$. We consider two cases: $S \in O$ and $S \in \{U, \bar{U}\}$. In the former case, let C be the coordinating strategy for S and let $k = k(1/2)$ be the number of miscoordination steps before C coordinates with probability $1/2$. For every j such that $S = S_j$, there is a positive probability of at least $1/2 \cdot 2^{-\lceil k/c_0 \rceil}$ that U will play C till there are k miscoordinations, and if that happens then with probability at least half it continue to achieve coordination with S. Since there are infinitely many j's such that $S = S_j$ we have that unless U gets stuck (in which case it has already coordinated) it will coordinate with S with probability 1 in a finite number of steps. (Note that we have greater control on the number of steps for coordination in this case - which depends on the frequency of S in the enumeration of O and the parameters c_0 and k).

In the case $S = U$ or $\bar{U}$, we have that in the beginning of a stage there is a chance of $1/16$ that U will play constant 0 until $2c_0$ miscoordinations, and a chance of at least $1/8$ that before these $2c_0$ miscoordinations are reached S will start a new stage in which it plays 0 as well. Therefore there is a positive probability that both U and S play the same constant at the beginning of a stage and this leads to coordination.

We note that U is uniformly coordinateable with the strategy C that plays all 0s with $k(\epsilon) = 2\lceil c_0 \log_2(1/\epsilon) \rceil$.

Finally we note that if the sequence $C_j = C(S_j)$ can be enumerated efficiently and C_j can be computed efficiently, then U is computable efficiently. (In particular U does not need to determine how long C will take to coordinate). □

The positive results show that eventual coordination is possible under rather limited assumption. Take O to be the set of all Turing computable (uniformly) coordinateable strategies is a countable set, and let U be the universal strategy constructed above. Note that O itself does not need to be common knowledge, as U is universal for any subset of O as well. In addition, U switches between actively trying to learn and passively playing a constant sequence. If we choose the constant sequence to replicate a strategy from O, it would allow any universal strategy for O to "learn" U and eventually coordinate with it.

7. Concluding Remarks

We model challenges in communication as a coordination game. The repeated coordination game gives the simplest instantiation of such a setting and allows us to conveys insights in a simple form. Using this framework we ask whether we can attain universality; can two players with limited prior knowledge of each other learn to meaningfully communicate by interacting repeatedly? We show we can exploit initial asymmetry in the player's role to achieve universality: one player will be assigned a "passive" role, and the other player will "actively learn". When players are ex-ante symmetric and labels are not informative universality cannot be achieved without randomization. Randomization allows the players to attempt to emulate initial asymmetry, which they can attempt repeatedly until reaching coordination. The players adopt strategies that are initially randomized to generate asymmetry, and deterministically leverage that asymmetry to reach full coordination.

From the communication point of view, this simple setting highlights the asymmetric role played by different players in solutions provided in previous works and explains why this asymmetry was essential. A symmetric solution can be achieved under limited assumptions on the knowledge of players, using computational limitations as a natural restriction. Since randomization is necessary the players can guarantee they will eventually establish an agreed encoding rule, but the encoding rule must be randomly selected.

References

Bergemann, Dirk and Stephen Morris, "Robust mechanism design," *Econometrica*, 2005, *73* (6), 1771–1813.

Bhaskar, V, "Egalitarianism and efficiency in repeated symmetric games," *Games and Economic Behavior*, 2000, *32* (2), 247–262.

Blume, Andreas, "Coordination and learning with a partial language," *Journal of Economic Theory*, 2000, *95* (1), 1–36.

———, **Douglas V DeJong, Yong-Gwan Kim, and Geoffrey B Sprinkle**, "Experimental evidence on the evolution of meaning of messages in sender-receiver games," *The American Economic Review*, 1998, *88* (5), 1323–1340.

Crawford, Vincent P and Hans Haller, "Learning how to cooperate: Optimal play in repeated coordination games," *Econometrica: Journal of the Econometric Society*, 1990, pp. 571–595.

———, **Uri Gneezy, and Yuval Rottenstreich**, "The power of focal points is limited: even minute payoff asymmetry may yield large coordination failures," *The American Economic Review*, 2008, *98* (4), 1443–1458.

Cremer, Jacques, Luis Garicano, and Andrea Prat, "Language and the Theory of the Firm," *The Quarterly Journal of Economics*, 2007, *122* (1), 373–407.

Dessein, Wouter, "Authority and communication in organizations," *The Review of Economic Studies*, 2002, *69* (4), 811–838.

Farrell, Joseph, "Cheap talk, coordination, and entry," *The RAND Journal of Economics*, 1987, pp. 34–39.

Gilboa, Itzhak and Dov Samet, "Bounded versus unbounded rationality: The tyranny of the weak," *Games and Economic Behavior*, 1989, *1* (3), 213–221.

Goldreich, Oded, Brendan Juba, and Madhu Sudan, "A theory of goal-oriented communication," *Journal of the ACM (JACM)*, 2012, *59* (2), 8.

Juba, Brendan and Madhu Sudan, "Universal semantic communication I," in "Proceedings of the 40th annual ACM symposium on Theory of computing" ACM 2008, pp. 123–132.

——— **and** ———, "Efficient Semantic Communication via Compatible Beliefs," in "Proceedings of the Second Symposium on Innovations in Computer Science-ICS" Institute for Computer Science, Tsinghua University 2011, pp. 7–9.

Kalai, Ehud, "Bounded rationality and strategic complexity in repeated games," *Game theory and applications*, 1990, pp. 131–157.

——— **and Ehud Lehrer**, "Rational learning leads to Nash equilibrium," *Econometrica: Journal of the Econometric Society*, 1993, pp. 1019–1045.

Knoblauch, Vicki, "Computable Strategies for Repeated Prisoner's Dilemma," *Games and Economic Behavior*, 1994, *7* (3), 381–389.

Kuzmics, Christoph, Thomas Palfrey, and Brian W Rogers, "Symmetric play in repeated allocation games," *Journal of Economic Theory*, 2014, *154*, 25–67.

Mailath, George J and Larry Samuelson, *Repeated games and reputations*, Vol. 2, Oxford university press Oxford, 2006.

Marschak, Jacob, "Economics of language," *Behavioral Science*, 1965, *10* (2), 135–140.

Nachbar, John H, "Prediction, optimization, and learning in repeated games," *Econometrica: Journal of the Econometric Society*, 1997, pp. 275–309.

———, "Beliefs in repeated games," *Econometrica*, 2005, *73* (2), 459–480.

——— **and William R Zame**, "Non-computable strategies and discounted repeated games," *Economic theory*, 1996, *8* (1), 103–122.

Rubinstein, Ariel, *Economics and language: Five essays*, Cambridge University Press, 2000.

Sandroni, Alvaro, "Reciprocity and cooperation in repeated coordination games: the principled-player approach," *Games and Economic Behavior*, 2000, *32* (2), 157–182.

Schelling, Thomas C, *The Strategy of Conflict*, Harvard university press, 1980.

Selten, Reinhard and Massimo Warglien, "The emergence of simple languages in an experimental coordination game," *Proceedings of the National Academy of Sciences*, 2007, *104* (18), 7361–7366.

Sipser, Michael, *Introduction to the Theory of Computation*, Vol. 2, Thomson Course Technology Boston, 2006.

Xandri, Juan Pablo, "Credible Reforms: A Robust Implementation Approach," 2012.