Universal Communication via Robust Coordination

Jacob D. Leshno* and Madhu Sudan†

November 20, 2013

Abstract

We consider the task of two players who wish to meaningfully communicate with each other, focusing on a simple and essential part of meaningful communication — the two players must coordinate on common interpretation of messages in order to be able to communicate. We present some basic definitions that capture the notion of the players *eventually* reaching a state of understanding: We formulate this problem as a repeated coordination game and ask whether the two players can guarantee that after a finite learning time they will coordinate in all periods forward. We ask for coordination in potentially varying environments under limited prior knowledge. Our results show that when there is some "grain of coordination" it can be leveraged to eventual coordination, but it is impossible to achieve coordination deterministically without some initial asymmetry. We give conditions under which randomization can be used to generate a "grain of coordination" and guarantee eventual coordination in a symmetric setting.

*Columbia Business School, Uris Hall, Columbia University, New York, NY 10027, USA. Part of the work was done when the author was a postdoctoral researcher at Microsoft Research New England. `jleshno@columbia.edu`.

†Microsoft Research New England, One Memorial Drive, Cambridge, MA 02142, USA. `madhu@mit.edu`.

1 Introduction

Universality, as interpreted as a single solution to all problems, is a desirable feature in any setting. In the context of computing this leads to the concept of a universal computer — a single computer capable of executing any program. In the context of data compression, this led to the notion of universal compression algorithms (like Lempel-Ziv). In the context of communication, an analogous notion may be a single communication protocol that can enable a person/device to meaningfully communicate with any other person/device. Is it possible to create such a universal communication protocol?

We build upon the work of Goldreich et al. [7] (see also Juba and Sudan [9, 10]) who studied this question in a broad setting. Their work asks whether communicating agents can achieve goal (typically computing a function), and suggests that if goals are well defined and agents can sense progress then one can essentially achieve universality.

In this paper we use a stylized game-theoretic model to study under which condition universal meaningful communication can be achieved. We focus on coordination problem, which is an essential part of meaningful coordination (see Example 1 below). While our model is highly stylized, it captures many of the essential elements of Goldreich et al. in a very simple setting and allows us to describe out work in a self-contained manner. Our setup allows us to show the essential role of randomization or ex-ante asymmetry for universal communication.

1.1 Communication as coordination

Alice and Bob are two players who wish to communicate meaningfully with each other, despite the fact that they have minimal prior information about each other.For Alice and Bob to communicate meaningfully requires more than just the passing of a message, Alice and Bob must also agree on the interpretation of the message. To clarify the difficulty we first present an example:

Example 1. Alice and Bob are randomly selected contestants on a game show. In each round a prize is hidden behind one of two curtains, and the other curtain has non prize. In each round Bob has to choose one of the two curtains, say left or right. Both players get the prize if it is hidden behind the chosen curtain. Alice is shown which curtain has the prize but is given a limited means to communicate her information to Bob. She is shown two "random" pictures and she gets to determine the order in which they are presented to Bob, after seeing which curtain has the prize. Based on the order Bob gets to pick a curtain to open and both players learn whether they won the round, and a new round starts. Can Alice and Bob win all prizes from some stage onwards?

In the example Alice needs to communicate one bit of information to Bob, and she can choose one of two possible messages. There are two possible messages, if the pictures are X and Y Alice can send the message $\langle X, Y \rangle$ or the message $\langle Y, X \rangle$. We assume that Alice is trying to transmit her information to Bob, and will therefore send a different message depending on whether the prize is behind the left or right curtain. Bob is trying to receive the information from Alice, and will therefore choose a different curtain depending on the message from Alice. But there are two possible encodings: In the first $\langle X, Y \rangle$ indicates that the left curtain has the prize and in the second $\langle Y, X \rangle$ indicates that the left curtain has the prize.

For Alice and Bob to meaningfully communicate and win the prize the must coordinate on the same encoding. If Alice and Bob use the same encoding they will win the prize, if they use different encoding they will not win[1]. We capture this problem using game theoretic tools, modeling the problem as a repeated coordination game. Each round is a *coordination game* — where Alice and Bob have to agree on a binary decisions. Their ability to "win"

[1] If Alice and Bob could agree on a global encoding, each rounds encoding could be derived from this global encoding. Unfortunately, the global encoding must be general enough to explain how to order arbitrary pairs of pictures, since in each round the pair of pictures available to Alice is different.

corresponds to their being able to coordinate and agree on a “common language”. At the end of each round they observe the outcome of the round and continue to play the next round. The interests of Alice and Bob are aligned, but they face strategic uncertainty. If they manage to resolve the strategic uncertainty they will manage to win all the prizes to follow. Note that resolving the strategic uncertainty is harder than coordinating in one period, since even if Alice and Bob managed to coordinate in a given round, they still face a challenge in the next round when they face a new set of messages.

A more standard version of a repeated coordination game is given in the following example:

Example 2. Alice and Bob are two business persons who are often sent to the same city. They would like to have dinner together, but in every town there are two equally good restaurants. They simultaneously choose a restaurant for dinner, and get a higher payoff if they end up eating together, having chosen the same restaurant.

Both examples are game-theoretically equivalent, both are a repeated coordination game with two actions. Each round of the game in example 2 is the classical coordination game with two actions, where the players get a higher payoff in every round if they choose the same option. Each round of the game in example 1 is also a coordination game with two actions, where the possible actions are the choice of encoding.[2]

Example 1 illustrates how coordination is essential for meaningful communication. To win a round in example 1 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 on an encoding or a language. Coordination games are in essence the simplest of games to be considered in the context of communication, as coordination is always required for the receiver to be able to interpret the message in the same way the transmitter meant.

[2] In both examples the two options change from period to period, but in game theoretic terms it is always the same game, and wlog we can refer to the two options simply as ’0’ and ’1’.

Furthermore, in typical settings players are indifferent between the different encodings, as long as both sides are in agreement. We therefore focus on the coordination game as an essential part of meaningful communication.

Without any prior agreement it is by definition impossible for Alice and Bob to meaningfully communicate a one-shot message, or coordinate in a one-shot coordination game. However, we can hope that coordination can be reached when the agents interact repeatedly and use observations from previous rounds to resolve the strategic uncertainty and agree on a strategy to coordinate. As is typical in other "universal settings", we should not expect the universal solution to be optimal (the universal Turing machine leads to descriptions of algorithms that is larger by a constant additive factor than the optimal). In our setting we would allow more than "optimal" time till coordination is reached, but hope that coordination can nevertheless be guaranteed in a finite amount of time.

We ask whether Alice and Bob can guarantee **eventual coordination**: Alice and Bob repeatedly play a repeated coordination game, with limited prior knowledge about each other, and they would like to guarantee that after some finite number of rounds (in which both get feedback) Alice and Bob coordinate in every round. Each player, given their (limited) knowledge would like to adopt a **universal strategy** which will guarantee eventual coordination, and meaningful communication, with the other player. We ask whether a universal strategy requires initial asymmetry or randomization.

1.2 Main Results

In our model we simply assume that Alice and Bob are playing a repeated coordination game with a binary choice and perfect monitoring (i.e., players observe each other's actions).[3] Each

[3] Returning to the motivation of modeling communication, we note that our repeated games do not allow any communication between players apart from the their actions in the repeated games. Such communication will also require coordination on common interpretation, and will thus have to be modeled as another coordination game (as we illustrate in example 1)

player chooses a strategy in the repeated game, which is a mapping from game histories of play to (a distribution over) actions. Specifically, we allow the players to have a different state after different plays. Since we do not assume that the agents have correct prior information of each other, we cannot expect the agents to play a Nash equilibrium. Instead, we model the agent i's (coarse) information by a set O, representing that i knows that the strategy of player j is $S_j \in O$. For example, the set O may be the set of Turing computable strategies, or representing any other restriction on the strategy space that i can assume. We ask if given the prior knowledge that Bob is playing some strategy in the set O Alice has a strategy that will guarantee eventual coordination. If so, we will call such a strategy for a **universal strategy for** O. We ask whether a universal strategy exists, and whether it requires the players to use randomization.

We start by considering an asymmetric problem, from Alice's perspective. Suppose Alice knows that Bob plays a strategy in O and wishes to guarantee eventual coordination with him (note that this is the typical setting considered in [7]). We ask: Is there a universal strategy for Alice which will guarantee eventual coordination? We show that under two restrictions on the set O the answer is positive. The first restriction is that every strategy in O is **coordinateable**, meaning that for every strategy in O and after every history of play there exists *some* strategy which will reach eventual coordination with O. The second requirement is that the set O is **countable**, a requirement that will be satisfied if for example we require that every strategy in O is Turing computable. Under these assumptions Alice has a Universal strategy for O which guarantees coordination whenever Bob plays a strategy from O. We show that these assumptions are essentially necessary, in the sense that there exist knowledge sets O that do not satisfy these requirements for which Alice will not be able to guarantee coordination.

We then proceed to ask if the existence of the 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 knowledge, they both know that they both play strategies in O; can they guarantee eventual coordination while keeping this knowledge true? If the universal strategy for Alice in the previous section had been in O, then the answer would have been positive. Indeed if there exists any universal strategy for O that is itself in O, then this would yield a symmetric solution.

If O consists of a single strategy of always playing, say, the action 0, then the only strategy in O is universal for O. However, this universality relies on distinguishing one of the two actions as "special"[4], which requires some prior coordination.[5] To preclude this prior coordination via the labels of the actions we ask that O is **label neutral**. A class of strategies O is label neural if for every strategy S in O there is also a strategy which is the same as S if we switch the names of the two actions, that is we can describe the set O when the strategy names are not meaningful.

We show that when O is label neutral eventual coordination cannot be guaranteed in a symmetric deterministic setting. In fact, any deterministic universal strategy for O does not belong to O. 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 attain coordination: the asymmetric 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 and hold symmetric beliefs on the labels of the game then even repeated play cannot guarantee eventual coordination. If both players try to be active learner or both players are passive they will not reach eventual coordination. While the proof is extremely simple and obvious in hindsight, it we note that this proof highlights the advantage of the simpler setting we propose in this paper (in contrast to those in [7]). An immediate and important corollary is that while there exists

[4]In both Examples 1 and 2 there is no natural labeling of the actions.

[5]Under this knowledge st O the two agents are able to coordinate in a one shot game as well, which indicates prior coordination rather than learned coordination.

a universal strategy for the class of Turing computable and coordinateable strategies, any such universal strategy is itself not Turing computable.

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. The rough intuition is that if Alice and Bob had a correlated coin they could break the symmetry, and commonly agree who is an active learner and who is passive. Two independent coins are perfectly correlated with probability half, so independent randomization can also help break the symmetry. By using strategies that guarantee that once coordination is reached it's never lost, and by exploiting the fact that the players can keep attempting while coordination was not reached, we design universal strategies that achieve eventual coordination from initial randomness.

Finally, we remark that the symmetric universal strategy is a highly desirable outcome in that it gives the moral equivalent of a dominating strategy for players. Given a knowledge set O a player does not have to choose between playing a strategy in O or a strategy that is universal for O. Choosing any universal strategy will guarantee eventual coordination, whether the other player plays a strategy in O or chooses a (potentially different) universal strategy.

1.3 Related and prior work

This paper is a part of a large literature studying the economics of language, starting with [12]. This literature explores the connections between the structure of language and decision making [13], structure of the firm and language [6] and experimentally investigates how subjects learn to communicate [3, 15]. Our paper contributes to this literature by showing that language can be learned from minimal structure.

We have already mentioned the work of Goldreich et al. [7] (see also Juba and Sudan [9, 10]). While our work is directly inspired by theirs, we feel the simplicity of the game-theoretic setting offers many clarifying insights. Goldreich et al. studied communication with complexity-theoretic goals and this prevented a clear separation of restrictions due to computational reasons specific to their goals from restrictions due to other reasons. In our simplified setting we manage to clarify the restrictions required for universality results.

Schelling's seminal work [14] 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. [4] 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.

There is extensive literature in game theory that has studied learning in repeated games. Several of these works focus on setting where players have limited information about each other. A seminal example is the work of Kalai and Lehrer [11] who study convergence to equilibrium through repeated play. They study general (and not just coordination) 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 consistent with Nash equilibrium. While this (and other results we mention next) have many similar ingredients to our study, we stress that Nash equilibrium is a very different solution concept that does not guarantee coordination. For instance two strategies that play 0/1 with probability 50% each independent of the history are in equilibrium, and would be far from our desired notion of eventual coordination.

Crawford and Haller [5] also study a similar setting to ours, where players try to coordinate when actions do not have ex-ante meaningful labels, but the action labels are fixed through the repeated game. They analyze the optimal strategies for quickest convergence to coordination. In terms of our model they make strong knowledge assumptions, they assume that once players coordinate in a single period they can maintain coordination forever by keep to play the same action. In our setting that successful coordination in a single period does not necessarily allow the players to maintain coordination. In example 2 coordination on a restaurant in Denver will not inform the players which restaurant they should choose in Salt Lake City. [2] shows that if agents have stronger knowledge of each other they can learn to coordinate more efficiently.

Hart and Mas-Colell [8] study a related but different question. In their setting players do not know payoffs in advanced. The players can agree on strategies, but they ask the strategies to be uncoupled, meaning that each player's strategy is allowed to depend only on his payoff (and not the other player's payoff). They show that no natural uncoupled strategy converges to Nash equilibrium.

Our solution concept deviates from the Nash equilibrium assumptions of common knowledge, and instead asks to adopt strategies that will work well even under a "worst-case" assumption. Bergemann and Morris [1] 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 [17] is an example of a paper that uses a similar solution concept to study reputation building.

Organization of this paper. In Section 2 we introduce our model formally and present our formal definition of eventual coordination. In Section 3 we consider the asymmetric setting and present some positive results. In Section 4 we consider the symmetric setting and present our negative result for deterministic strategies. In Section 5 we consider the symmetric setting and present positive results using randomness. Some concluding thoughts are presented in Section 6.

2 Model and Definitions

In Section 1 we argued that coordination is an essential part of meaningful communication. In this section we present our formal framework, which models the problem as a repeated coordination game. We proceed to define the game and our goals.

We model the interaction between the agents as a repeated coordination game with two players. We use $i \in \{1,2\}$ to denote a player and $t \geq 0$ to denote a period. In each period $t \geq 0$ player i chooses an action $a_i^t \in A_i = \{0,1\}$.[6] The action profile of period t is $a^t = (a_1^t, a_2^t) \in A = A_1 \times A_2$. The players coordinate in period t if $a^t \in \mathcal{C} = \{(0,0),(1,1)\}$. Both players receive a payoff of +1 in periods where they coordinate and a payoff of 0 in periods where they do not coordinate. Thus there is no conflict of interest between the two players, but players face strategic uncertainty when they have incomplete knowledge of the other player's strategy. The goal of both players is to guarantee eventual coordination, formally defined below, which will guarantee that their long run average payoff is equal to +1.

After each period both players observe the actions of both players and continue to play the next period. Their actions may thus depend on

At the end of period t both players observe the action of the other player[7], and therefor learn a^t. The history at (the beginning of) period t, denoted h^t, is defined to be the sequence of action profiles up to period $t-1$, i.e., $h^t = (a^0, a^1, \ldots, a^{t-1})$. We use $H = (A_1 \times A_2)^*$ to denote the set of all possible histories.

A strategy $S_i : H \to \Delta(A_i)$ for player i specifies the (random) action the player chooses given the observed information. With slight abuse of notation, we refer to deterministic strategy S_i as a mapping $S_i : H \to A_i$. We can equivalently describe the strategies as state

[6]In Examples 1,2 the set of actions changed from period to period. For ease of exposition we use the same labels for actions of different rounds, while keeping in mind that our analysis must treat the names of actions as arbitrary labels.

[7]In example 1 both players learn whether they won the prize after each period, and therefore they learn whether the other player used the same encoding.

dependent strategies; In this view, a strategy for player i is described by a pair of functions:

$$S_i : \Sigma_i \to \Delta(A_i)$$

$$\Phi : \Sigma_i \times A \to \Sigma_i$$

where Σ_i is an abstract (not necessarily finite) set denoting the state space of S_i. The strategy takes a state $\sigma_i^t \in \Sigma_i$ at the beginning of period t, determines the current action $a_i^t = S_i(\sigma_i^t)$ at time t, and after observing the joint action profile a^t determines the next state $\sigma_i^{t+1} = \Phi(\sigma_i^t, a^t)$. This extends the previous definition by setting $\Sigma_i = H$ 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.

Let $\Sigma = \prod_i \Sigma_i$ and $A = \prod_i A_i$. We denote by $S_i(\sigma_i)$ the strategy S_i starting from the state σ_i. We denote by ϕ the state corresponding to the initial state, which corresponds to the empty history, and denote by $S_i(\phi)$ or simply S_i the strategy S_i starting from the initial state. A strategy profile $S = (S_1(\sigma_1), S_2(\sigma_2))$ defines a mapping $S : \Sigma \to \Delta(A \times \Sigma)$ and initial state $(\sigma_1, \sigma_2) \in \Sigma$, and therefore it defines a probability distribution $\mathbb{P}_S$ over H describing the probability of future paths of play. A deterministic strategy profile S and initial states σ_1, σ_2 define a deterministic sequence of play $h \in H$.

The next definitions captures the goal of the players, formalizing our notion of achieving coordination through repeated interaction.

Definition 1. Two strategies $S_1(\sigma_1), S_2(\sigma_2)$ **eventually coordinate** if the induced path of play includes only a finite number of non-coordination periods with probability 1. That is,

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

The players face no conflict of interest, but they face strategic uncertainty. If the players manage to resolve the strategic uncertainty after a finite number of periods they will be able to achieve the coordination payoff from all periods onwards. Note that we allow σ_1 and σ_2 to correspond to inconsistent histories.

We model the initial knowledge of each player as a set O of possible opponent's strategies. A element of O is a pair $(S_j, \sigma_j) \in O$ where S_j is a potential strategy and σ_j is its initial state. If $(S_j, \sigma_j) \notin O$ each of the players is certain that the opponent is not playing (S_j, σ_j). With slight abuse of notation we write $S \in O$ for $(S, \phi) \in O$. We ask whether the player can guarantee eventual coordination given its knowledge set O.

Definition 2. Player i who has knowledge O can **guarantee coordination** if there exists a strategy U_i for player i such that for every $(S_j, \sigma_j) \in O$ the strategies $U_i, S_j(\sigma_j)$ eventually coordinate. We refer to such U_i as a **universal strategy for O.**

Note that if U_i is universal for O and $O' \subset O$ then U_i is universal for O' as well. Thus designing universal strategies for larger knowledge sets is more challenging. Often our sets O will include strategies with all their possible states. We will try to let O be as large as possible, and exclude strategies by making minimal assumptions which are necessary for universality.

3 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 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}. The example above could be generalized to any other "inherently randomized" strategy[8]. 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 at least allow coordination with some strategy. We stress that the quantifiers are "switched": Our requirement here is that for every strategy $S_j(\sigma_j) \in O$ there is some strategy $C = C(S_j, \sigma_j)$ that achieves eventual coordination with $S_j(\sigma_j)$. In contrast the condition for universality is that there should exist some strategy U_i such that U_i achieves eventual coordination with $S_j(\sigma_j)$ for every $S_j(\sigma_J) \in O$.

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[9], else it mixes 50-50 forever.

$S^p_{\text{password1try}}(\phi)$ 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. Coordinating with a strategy $S^p_{\text{password1try}}$ requires a particular play of the first ten periods, which is incompatible with the required play of $S^q_{\text{password1try}}$. 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.

[8]That is, any strategy such that plays both actions with probability $> \varepsilon > 0$ after every history.

[9]Recall that the names of actions are meaningless. "Playing 0 forever" simply means playing a fixed sequence of actions.

Definition 3. We say that a strategy $S_j(\sigma_j) \in O$ is **coordinateable**, if, for every possible state σ'_j that can be reached from σ_j by S_j there exists a strategy $C = C^{(S_j,\sigma_j)}$ such that $C(\phi)$ [10] eventually coordinates with $S_j(\sigma_j)$. We say that O is **coordinateable** if every $S_j(\sigma_j) \in O$ is coordinateable.

This brings us to our first requirement:

Axiom 1. *The set O 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 (or some other deterministic sequence that the players can coordinate on). Else, it moves to p_2 for the next ten periods, etc.

$S_{password}$ is coordinateable, because for every history there is a sequence of play such that within 20 periods $S_{password}$ will play only 0. A strategy that knows the relevant p_k will quickly coordinate with $S_{password}$, but for a strategy that does not know $\{p_k\}$ it will not be necessarily possible to coordinate. If O includes all possible $S_{password}$ i.e., for every possible sequence of $\{p_k\}_k$ 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:

[10] i.e., C with its default initial state .

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 the opposing agent has limited computation power, limiting him to computable strategies[11], or more restrictively, to efficiently computable strategies. Since the set of all computable strategies is countable this restriction allows such natural classes of strategies.

The following theorem shows that universality is possible under the two assumptions. In the theorem we restrict attention to deterministic strategies[12], showing that even without randomization player i can have a deterministic universal strategy.

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

Remark. The proof below can be derived from the work of [7], and included here for completeness.

Proof. We build U as follows. Let $\{S_j(\sigma_j) : j \in \mathbb{N}\}$ 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 $S_j(\sigma_j)$, then skip to stage $j+1$, else let ρ_j denote the state of strategy $S_j(\sigma_j)$ 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 $S_j(\sigma_j)$ is coordinateable) and let n_j be an upper bound on the number of non-coordination periods before $C_j(\phi)$ 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 $j+1$.

[11] A strategy is computable if the function describing the strategy can be generated by a Turing machine (for further details see [16]).

[12] The treatment of randomized strategies is delayed to Section 5.

To verify that U is universal, suppose it is playing against $S_j(\sigma_j)$. 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 $S_j(\sigma_j)$. $\square$

The following corollary is immediate, using the fact that the set of computable strategies is countable.

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

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.

4 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.

The previous section shows that if the players have previously coordinated on asymmetric roles they can leverage this to achieve eventual coordination.

In this section we ask whether the players can achieve eventual coordination without prior coordination. We require that the players will have symmetric roles, 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.

Example 6 shows a trivial knowledge set which makes symmetric universal coordination possible. But the set O hardly satisfies our initial goal to enable eventual coordination under minimal knowledge. Coordination in the above example is not learned through repeated play; rather, play is coordinated from the first period. This coordination is facilitated by giving the action 0 a special role not available to the action 1. Indeed a knowledge set that gives a special role to one of the two actions requires some "prior" agreement. This motivates us to study knowledge sets that are neutral between the two actions 0 and 1. We formalize this concept next.

For $a \in \{0, 1\}$ we denote by $\overline{a} \triangleq 1 - a$ the label-switched action. For an action profile $a^t = (a_1^t, a_2^t)$, we let $\overline{a^t} = (\overline{a_1^t}, \overline{a_2^t})$ denote its label-switched profile. For a history $h = (a^0, a^1, \ldots, a^t)$ we let $\overline{h}$ denote its label-switched history, where $\overline{h} = (\overline{a^0}, \overline{a^1}, \ldots, \overline{a^t})$. 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 with the labels of the actions switched, both on the history it observes and on the actions it outputs.

Note that for the empty history ϕ, $\overline{\phi} = \phi$ and so $\overline{S}(\phi) = \overline{S(\phi)}$ - in other words, the first action of $\overline{S}$ is the complement of the first action of S.

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

The assumption that O is label neutral simply means that action names are not meaningful on their own, and therefore the knowledge set O should remain the same if we switch the names of the two actions. 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 ϕ, 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 can not be universal. $\square$

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. In asymmetric setting explored in the previous section, the players reached eventual coordination by designating one player as an "active learner" and the other as "passive". The active learner played a universal strategy that eventually figured out how to coordinate with the other player's strategy. For this universal strategy to work, the other strategy needs to "stay put" in a sense, allowing the other player to learn it. What would happen if the other player also attempted to learn? The theorem above shows that unless we can preclude that both players are "active learners" (both trying to learn each other), then 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 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. 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.

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.

5 Symmetric Coordination: randomized possibility results

In the previous section we saw that if the players have some "grain of coordination" which allows an asymmetry between the player's roles it can be leveraged to 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'.*

We prove the result by showing that randomization can create the asymmetry which enabled our positive results in Section 3. 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 3. 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. For simplicity we 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 $\{S_j(\sigma_j) | j \in \mathbb{N}\}$ 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 $S_j(\sigma_j)$ and tries to coordinate with it as follows. Let h_j be the history at the beginning of stage j. If the probability of h_j is zero then we terminate the stage. Else let ρ_j be the state of S_j starting with state σ_j and with history h_j. Let $C = C(S_j, \rho_j)$ be the strategy that coordinates with $S_j(\rho_j)$ and let n_j be the minimal integer such that the probability that $C(\phi)$ 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 $\{S_{j'}(\sigma_{j'})|j' \leq j\}$ 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 $S_j(\sigma_j) \in O'$. We consider two cases: $S_j(\sigma_j) \in O$ and $S_j(\sigma_j) \in \{U, \bar{U}\}$. In the former case, we have that unless U gets stuck in some stage, it tries to coordinate with $S_j(\sigma_j)$ infinitely often and each time it has a $1/2$ chance of achieving coordination conditioned on the past. 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 $S = U$ or $S = \bar{U}$ then either coordination is reached, or it happens infinitely often that one of them initiates stage k' while the other strategy 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. 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 $k < k'$ is reached, the probability of reaching coordination is at least $1/16$, leading to eventual coordination with probability one in finite number of periods. □

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 $C(S_j, \tau_j)$ 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.

5.1 Untraceable states and uniform coordination

Our model allows players to have infinite recall, allowing them to calculate the opponents current state given an hypothesized strategy and history of play. The universal strategies we constructed so far required this ability to calculate the appropriate strategy which will coordinate with the opponent at his current state. We follow to strengthen our definition of coordinateability, which will allow us to relax this requirement.

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 $C(\phi)$ 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 requires the universal strategy to eventually coordinate only with uniformly coordinateable strategies, it requires U itself to be universally coordinateable as well.

Remark 1. If the set of strategies $\{C(S)|S \in O\}$ 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 for simplicity we assume O is label neutral.

We start by describing the $O' = O \cup \{U, \bar{U}\}$ universal strategy U: Let $c_0 \geq 1$ be some fixed constant. Let $\{S_j|j \in \mathbb{N}\}$ be an enumeration of strategies in O in which every element of O appears infinitely often (note that uniformity allows us to ignore the states of S_j). 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 $C(\phi)$ in phases of length c_0. After c_0 miscoordination steps, U tosses a coin and with probability $1/2$ it aborts this stage and moves to stage $j+1$, and with probability $1/2$ it continues the stage moving to the next phase.

To prove universality of U for O' fix a strategy $S(\sigma) \in O'$. We consider two cases: $S(\sigma) \in O$ 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 with positive probability 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 determined enumerated efficiently and C_j can be computed efficiently, then U is computable as efficiently. (In particular U

does not need to determine how long C will take to coordinate). $\square$

6 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. 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, and how symmetry can be achieved by randomization.

Finally, our methodology differs from standard equilibrium concepts, instead focusing on a "worst-case" solution concepts. We find this approach particularly fitting for our questions. This framework allows us to investigate the necessary requirements of coordination, and show that under minimal requirements coordination can be guaranteed. The positive results suggest that such concepts may indeed by useful. In settings where there is much uncertainty and players are patient, our universal strategies provide a solution which allows the players to guarantee coordination in the long run.

References

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

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

[3] Andreas Blume, 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*, 88(5):1323–1340, 1998.

[4] Vincent P Crawford, 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*, 98(4):1443–1458, 2008.

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

[6] Jacques Cremer, Luis Garicano, and Andrea Prat. Language and the theory of the firm. *The Quarterly Journal of Economics*, 122(1):373–407, 2007.

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

[8] Sergiu Hart and Andreu Mas-Colell. Uncoupled dynamics do not lead to nash equilibrium. *The American Economic Review*, 93(5):1830–1836, 2003.

[9] Brendan Juba and Madhu Sudan. Universal semantic communication i. In *Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 123–132. ACM, 2008.

[10] Brendan Juba and Madhu Sudan. Efficient semantic communication via compatible beliefs. In *Proceedings of the Second Symposium on Innovations in Computer Science-ICS*, pages 7–9. Institute for Computer Science, Tsinghua University, 2011.

[11] Ehud Kalai and Ehud Lehrer. Rational learning leads to nash equilibrium. *Econometrica: Journal of the Econometric Society*, pages 1019–1045, 1993.

[12] Jacob Marschak. Economics of language. *Behavioral Science*, 10(2):135–140, 1965.

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

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

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

[16] Michael Sipser. *Introduction to the Theory of Computation*, volume 2. Thomson Course Technology Boston, 2006.

[17] Juan Pablo Xandri. Credible reforms: A robust implementation approach. 2012.