

Self-stabilizing uncoupled dynamics^{*}

Aaron D. Jaggard¹, Neil Lutz², Michael Schapira³, and Rebecca N. Wright⁴

¹ U.S. Naval Research Laboratory, Washington, DC 20375, USA.
aaron.jaggard@nrl.navy.mil

² Rutgers University, Piscataway, NJ 08854, USA. njlutz@cs.rutgers.edu

³ Hebrew University of Jerusalem, Jerusalem 91904, Israel. schapiram@huji.ac.il

⁴ Rutgers University, Piscataway, NJ 08854, USA. rebecca.wright@rutgers.edu

Abstract. Dynamics in a distributed system are self-stabilizing if they are guaranteed to reach a stable state regardless of how the system is initialized. Game dynamics are uncoupled if each player’s behavior is independent of the other players’ preferences. Recognizing an equilibrium in this setting is a distributed computational task. Self-stabilizing uncoupled dynamics, then, have both resilience to arbitrary initial states and distribution of knowledge. We study these dynamics by analyzing their behavior in a bounded-recall synchronous environment. We determine, for every “size” of game, the minimum number of periods of play that stochastic (randomized) players must recall in order for uncoupled dynamics to be self-stabilizing. We also do this for the special case when the game is guaranteed to have unique best replies. For deterministic players, we demonstrate two self-stabilizing uncoupled protocols. One applies to all games and uses three steps of recall. The other uses two steps of recall and applies to games where each player has at least four available actions. For uncoupled deterministic players, we prove that a single step of recall is insufficient to achieve self-stabilization, regardless of the number of available actions.

1 Introduction

Self-stabilization is a failure-resilience property that is central to distributed computing theory and is the subject of extensive research (see, e.g., [3] for a survey). It is characterized by the ability of a distributed system to reach a stable state from every initial state. Dynamic interaction between strategic agents is a central research topic in game theory (see, e.g., [4,11]). One area of interest is *uncoupled dynamics*, in which each player’s strategy is independent of the other players’ payoffs [9]. Here, we bring together these two research areas and study of *self-stabilizing uncoupled dynamics* within the broader research agenda of *distributed computing with adaptive heuristics* [10]. The same questions we answer here can be asked for a broad variety of dynamics and notions of convergence and equilibria. These directions, as well as a conjecture, are discussed in Section 5.

We focus our investigation on a bounded-recall, synchronous setting. We consider self-stabilization in a multi-agent distributed system in which, at each

^{*} Research supported in part by NSF Grant No. CCF-1101690 and by ONR.

timestep, the agents act as strategic players in a game, simultaneously selecting *actions* from their respective finite action sets to form an *action profile*. The space of action profiles is relevant throughout this work, and we refer to its size as the *size* of the game. We study the effects of *bounded recall*, in which the state of this system at any time consists of the r most recent action profiles, for some finite r . The stable states in r -recall systems necessarily have the same action profile in r consecutive time steps. In our context, we want stable states that are robust to players acting selfishly—i.e., those where the repeated action profile is an equilibrium of the stage game. In this paper, we consider pure Nash equilibria (PNE). Thus, in our setting, dynamics self-stabilize for a given game if, from every starting state, players are guaranteed to converge to a PNE. For games without PNE, dynamics cannot self-stabilize in this sense. Throughout this paper, we say that particular dynamics *succeed* on a class of games if they self-stabilize for games in that class whenever a PNE exists.

Traditional study of convergence to equilibria in game dynamics makes various assumptions about the “reasonableness” of players’ behavior, restricting them to always play the game in ways that are somehow consistent with their self-interest given their current knowledge. In contrast to these behavioral restrictions on the players, uncoupledness is an informational restriction, in that the players have no knowledge of each other’s payoffs. In this situation, no individual player can recognize a PNE, so finding an equilibrium is a truly distributed task.

If uncoupledness is the only restriction on the dynamics, then the players can find a PNE through a straightforward exhaustive search. However, this changes when players’ abilities to remember past actions is restricted. In a continuous-time setting, Hart and Mas-Colell [7] showed that deterministic uncoupled dynamics fail to reach a stable state for some games that have PNE if the dynamics must be *historyless*, i.e., if the state space of the system is identical to the action profile space of the game. This suggests the central question that we address:

On a given class of games, how much recall do uncoupled players need in order to self-stabilize whenever a PNE exists? That is, when are there successful k -recall dynamics?

This question was answered in part by Hart and Mas-Colell [8], who showed that in a discrete-time setting, even when players are allowed randomness, no historyless uncoupled dynamics succeed on all two-player games where each player has three actions. Moreover, they showed that even for *generic* games (where at every action profile each player has a unique best response), no historyless uncoupled dynamics succeed on games with three three-action players. They also gave positive results, proving that there are historyless uncoupled dynamics that succeed on all two-player generic games, and that if the players have 2-recall (i.e., they are allowed to see the two most recent action profiles), then over every action profile space there are stochastic uncoupled dynamics that succeed on all games.

Our results. We show in Section 3 that there exist historyless uncoupled dynamics that succeed on all two-player games with a two-action player and on

all three-player generic games with a two-action player (Theorems 4 and 9). In both cases, we prove that these results are tight, in that they do not hold for any larger size of game (Theorems 5 and 11). Combined with the results of Hart and Mas-Colell [8], this provides a complete characterization of the exact minimum recall needed, for any action profile space, for uncoupled dynamics to succeed on all games over that space and on generic games over that space. In Section 4, turning to deterministic dynamics, we demonstrate 3-recall deterministic uncoupled dynamics that succeed on all games (Theorem 14) and 2-recall deterministic uncoupled dynamics that succeed on all games in which every player has at least four actions (Theorem 15). We also prove for every action profile space that no historyless deterministic uncoupled dynamics succeed on all games over that space (Theorem 16). Some proofs are omitted from the proceedings version of this paper. A longer version with all proofs included can be found at <http://arxiv.org/abs/1403.5791>.

Related work. There are rich connections between distributed computing and game theory, some of which are surveyed by Halpern [5]. Jaggard, Schapira, and Wright [10] investigated convergence to pure Nash equilibria by game dynamics in asynchronous distributed systems. Most closely related to our specific setting, Hart and Mas-Colell introduced the concept of uncoupled game dynamics [7]. In addition to the results mentioned above, they also addressed convergence to mixed Nash equilibria by bounded-recall uncoupled dynamics [8]. Babichenko investigated the situation when the uncoupled players are finite-state automata, as well as *completely uncoupled dynamics*, in which each player can see only the history of its own actions and payoffs [1,2]. Young [13] and Pradeliski and Young [12] gave completely uncoupled dynamics that achieve an equilibrium in a high proportion of steps but do not necessarily converge. Hart and Mansour [6] analyzed the time to convergence for uncoupled dynamics.

2 Preliminaries

We begin with definitions of the concepts used in the paper.

Games. Let $n \in \mathbb{N}$ and $(k_1, \dots, k_n) \in \mathbb{N}^n$, with $n \geq 2$ and each $k_i \geq 2$. A *game* of size (k_1, \dots, k_n) is a pair (A, U) , where $A = A_1 \times \dots \times A_n$ such that each $|A_i| = k_i$, and $U = (u_1, \dots, u_n)$ is an n -tuple of functions $u_i : A \rightarrow \mathbb{R}$. A_i and u_i are the *action set* and *utility function* of *player* i . $\Delta(A_i)$, the probability simplex over A_i , is player i 's set of *mixed actions*. When n is small, we may describe a game (A, U) as a k_1 -by-...-by- k_n game. Elements of A are the *(action) profiles* of the game, and A is called the *(action) profile space*. $\mathcal{U}(A)$ is the class of all U such that each u_i takes A_i as input, so $A \times \mathcal{U}(A)$ is the class of all games with profile space A . When A is clear from context, we often identify the game with the utility function vector U .

Let $U \in \mathcal{U}(A)$. For $i \in \{1, \dots, n\}$ and $a = (a_1, \dots, a_n) \in A$, we say that player i is *U -best-replying* at a if $u_i(a) \geq u_i((a_1, \dots, a'_i, \dots, a_n))$ for every $a'_i \in A_i$. We

define the set of U -best-replies for player i at a ,

$$BR_i^U(a) = \{a'_i \in A_i : i \text{ is } U\text{-best-replying at } (a_1, \dots, a'_i, \dots, a_n)\}.$$

We omit U from this notation when the game being played is clear from context. A profile $p \in A$ is a *pure Nash equilibrium*, abbreviated *PNE*, for U if every player $i \in \{1, \dots, n\}$ is best-replying at p . An action $a_i \in A_i$ is *weakly dominant* for player i if $a_i \in BR_i(x)$ for every $x \in A$; it is *strictly dominant* for player i if $BR_i(x) = \{a_i\}$ for every $x \in A$.

A game $(A, U) \in A \times \mathcal{U}(A)$ is *generic* if every player's best-replies are unique, i.e., if for every $a \in A$ and $i \in \{1, \dots, n\}$, $|BR_i^U(a)|=1$. For generic games (A, U) we may abuse notation slightly by using $BR_i^U(a)$ to refer to this set's unique element. $A \times \mathcal{G}(A)$ is the class of all generic games on A .

Dynamics. We now consider the repeated play of a game. Let the profile at timestep $t \in \mathbb{Z}$ be $a^{(t)} = (a_1^{(t)}, \dots, a_n^{(t)}) \in A$. The *stage game* $(A, U) \in A \times \mathcal{U}(A)$ is then played: each player i simultaneously selects a new action $a_i^{(t+1)}$ by applying an r -recall *stationary strategy* $f_i^U : A^r \rightarrow \Delta(A_i)$, where $r \in \mathbb{N}$ and A^r is the Cartesian product of A with itself r times. A *deterministic* r -recall stationary strategy mapping ranges over A_i instead of $\Delta(A_i)$. The strategy f_i^U , which is *stationary* in the sense that it does not depend on t , will take as input $(a^{(t-r+1)}, \dots, a^{(t)})$, the r most recent profiles. We call this r -tuple the *state* at time t . The terms *1-recall* and *historyless* are interchangeable. A *strategy vector* is an n -tuple $f^U = (f_1^U, \dots, f_n^U)$, where each f_i^U is a strategy for player i . $\mathcal{F}(A)$ will denote the set of all strategy vectors for A .

A *strategy mapping* for A is a mapping $f : \mathcal{U}(A) \rightarrow \mathcal{F}(A)$ that assigns to each U a strategy vector f^U . A strategy mapping f is *uncoupled* if the strategy it assigns each player depends only on that player's utility function and not, e.g., on the other players' payoffs. That is, there are mappings f_1, \dots, f_n where each f_i maps utility functions on A to strategies for A , such that $f_i(u_i) \equiv f_i^U$ for $i = 1, \dots, n$. If f_i^U is stationary, deterministic, or r -recall for $i = 1, \dots, n$, then f^U is also. If every f^U has any of those properties, then f does also.

Now let $x = (x^{(1)}, \dots, x^{(r)}) \in A^r$, and let f^U be an r -recall strategy vector. For $T \geq r$, a *partial f^U -run* for T steps starting from x is a tuple of profiles $(a^{(1)}, \dots, a^{(T+r)}) \in A^{T+r}$ such that $x = (a^{(1)}, \dots, a^{(r)})$ and for every $r < t \leq T+r$,

$$\Pr \left(f^U(a^{(t-r)}, \dots, a^{(t-1)}) = a^{(t)} \right) > 0.$$

An f^U -run is an infinite sequence of profiles $a^{(1)}, a^{(2)}, \dots$ such that every finite prefix is a partial f^U -run. We say that $y \in A^r$ is f^U -reachable from $x \in A^r$ if there exist a $T \in \mathbb{N}$ and a partial f^U -run $(a^{(1)}, \dots, a^{(T+r)})$ such that $x = (a^{(1)}, \dots, a^{(r)})$ and $y = (a^{(T)}, \dots, a^{(T+r)})$. The state x is an f^U -absorbing state if for every f^U -run $a^{(1)}, a^{(2)}, \dots$ beginning from x , $(a^{(t+1)}, \dots, a^{(t+r)}) = x$ for every $t \in \mathbb{N}$. Notice that any f^U -absorbing state $x = (a^{(1)}, \dots, a^{(r)})$ must have $a^{(1)} = \dots = a^{(r)}$. We omit the strategy vector from this notation when it is clear from context. The *game dynamics* of f consist of all pairs (U, R) such that R is an f^U -run.

Convergence. A sequence of profiles $a^{(1)}, a^{(2)}, \dots$ *converges* to a profile a if there some $T \in \mathbb{N}$ such that $a^{(t)} = a$ for every $t \geq T$. If from every $x \in A^r$, some f^U -absorbing PNE is f^U -reachable, then f *self-stabilizes* on game (A, U) . We say that f *succeeds* on a game U if f self-stabilizes on (A, U) or if (A, U) has no PNE. Let $\mathcal{C}(A)$ be a class of games on A . If f succeeds on every game $(A, U) \in A \times \mathcal{C}(A)$, then f *succeeds* on $\mathcal{C}(A)$.

Let $A = A_1 \times \dots \times A_n$ and $B = B_1 \times \dots \times B_n$ be profile spaces of the same size, in the sense that there is some permutation π on $\{1, \dots, n\}$ such that $(|A_1|, \dots, |A_n|) = (|B_{\pi(1)}|, \dots, |B_{\pi(n)}|)$. Then we write $A \simeq B$. If f succeeds on $\mathcal{C}(A)$, then there is a strategy mapping derived from f that succeeds on $\mathcal{C}(B)$, simply by rearranging the players and bijectively mapping actions in each A_i to actions in $B_{\pi(i)}$. This new strategy mapping retains any properties of f that are of interest here (uncoupledness, r -recall, stationarity, and determinism). For this reason we define

$$\mathcal{C}(|A_1|, \dots, |A_n|) = \bigcup_{B \simeq A} \mathcal{C}(B),$$

and we say that f succeeds on $\mathcal{C}(|A_1|, \dots, |A_n|)$ if f succeeds on $\mathcal{C}(B)$ for some $B \simeq A$. For example, “ f succeeds on $\mathcal{G}(2, 3)$ ” means “ f self-stabilizes on every generic 2-by-3 game with a PNE (up to renaming of actions).”

3 Stochastic uncoupled dynamics

In this section, we determine, for every profile space A , the minimum $r \in \mathbb{N}$ such that an uncoupled r -recall stationary strategy mapping exists that succeeds on all games $(A, U) \in A \times \mathcal{U}(A)$ or all generic games $(A, U) \in A \times \mathcal{G}(A)$. Hart and Mas-Colell [8] proved that 2-recall is sufficient to succeed on all games, 1-recall is sufficient to succeed on generic two-player games, and that 1-recall is not sufficient to succeed on all games, or even all generic games. We state these results in the present setting.

Theorem 1 (Hart and Mas-Colell [8]). *For any profile space A , there exists an uncoupled 2-recall stationary strategy mapping that succeeds on all games (A, U) .*

Theorem 2 (Hart and Mas-Colell [8]). *There is no uncoupled historyless stationary strategy mapping that succeeds on all 3-by-3 games, or on all 3-by-3-by-3 generic games.*

Theorem 3 (Hart and Mas-Colell [8]). *For any two-player profile space A , there is an uncoupled historyless stationary strategy mapping that succeeds on all generic games (A, U) .*

We now describe the strategy mapping given in the proof of Theorem 3. Notice that for a historyless stationary strategy mapping, the state space is exactly the profile space, so the terms *state* and *profile* are interchangeable in this context.

Definition For any n -player profile space A , the *canonical* historyless uncoupled stationary strategy mapping for A is $h : \mathcal{U}(A) \rightarrow \mathcal{F}(A)$, defined as follows. Let $U = (u_1, \dots, u_n) \in \mathcal{U}(A)$. Then $h(U) = (h_1^U, \dots, h_n^U)$, where for $i \in \{1, \dots, n\}$, $h_i^U : A \rightarrow A_i$ is given by

$$\begin{aligned}\Pr(h^U(a_i) = a_i \mid a_i \in BR_i(a)) &= 1 \\ \Pr(h^U(a_i) = b_i \mid a_i \notin BR_i(a)) &= 1/k_i,\end{aligned}$$

for all $a_i, b_i \in A_i$. That is, if player i is already best replying, then it will continue to play the same action. Otherwise, i will play an action chosen uniformly at random from its action set.

In their proof of Theorem 2, Hart and Mas-Colell make the following observation.

Observation 1 (Hart and Mas-Colell [8]). *Suppose f is an uncoupled historyless stationary strategy mapping for profile space A and f succeeds on all generic games (A, U) . Then two conditions hold for every game (A, U) and $a = (a_1, \dots, a_n) \in A$. First, if player i is best-replying at a , then $\Pr(f_i^U(a) = a_i) = 1$. Second, if player i is not best-replying at a , then $\Pr(f_i^U(a) = a'_i) > 0$ for some $a'_i \in A_i \setminus \{a_i\}$.*

Informally, no player can move when it is best-replying, and each player must move w.p.p. whenever it is not best-replying. The first condition guarantees that every PNE is an absorbing state; the second guarantees that no non-PNE is an absorbing state. Implicit in the same proof is the fact that h is at least as “powerful” as any other historyless uncoupled strategy mapping.

Observation 2 (Hart and Mas-Colell [8]). *If any historyless uncoupled stationary strategy mapping succeeds on $\mathcal{U}(A)$ or on $\mathcal{G}(A)$, then h succeeds on that class.*

3.1 Stochastic dynamics for $\mathcal{U}(A)$

We now describe the profile spaces in which there are uncoupled historyless strategy mappings that succeed on every game, or equivalently (by Observation 2), the A for which h succeeds on $\mathcal{U}(A)$. A proof that h succeeds on 2-by- k games proceeds by simple case checking.

Theorem 4. *For every two-player profile space A in which one player has only two actions, h succeeds on all games (A, U) .*

It turns out that 2-by- k profile spaces are the only ones where h succeeds on all games.

Theorem 5. *Let A be a profile space. Unless A has only two players and one of those players has only two actions, no historyless uncoupled stationary strategy mapping succeeds on all games (A, U) .*

We give three lemmas that will be used in the proof of Theorem 5. Informally, Lemma 6 says that additional actions do not make a profile space any “easier” in this context; the players will need at least as much recall to succeed on all games in the larger space. The proof relies on a type of reduction in which the players take advantage of a strategy mapping for a larger game by “pretending” to play the larger game. Whenever player i plays k_i , all players guess randomly whether i would have played k_i or $k_i + 1$ in the larger game.

Lemma 6. *Let $n \geq 2$, $k_1, \dots, k_n \geq 2$, and $i \in \{1, \dots, n\}$. If h succeeds on $\mathcal{U}(k_1, \dots, k_i + 1, \dots, k_n)$, then h succeeds on $\mathcal{U}(k_1, \dots, k_i, \dots, k_n)$.*

Lemma 7 tells us that the same is true of adding players to the game. Its proof also uses a simple reduction. The players utilize the strategy mapping for the $(n + 1)$ -player game by behaving as if there is an additional player who never wishes to move. This preserves genericity, so the lemma also applies to the class of generic games.

Lemma 7. *Let $n \geq 2$ and $k_1, \dots, k_n, k_{n+1} \geq 2$. If h succeeds on $\mathcal{U}(k_1, \dots, k_n, k_{n+1})$, then h succeeds on $\mathcal{U}(k_1, \dots, k_i, \dots, k_n)$. The same is true if we replace \mathcal{U} with \mathcal{G} .*

Finally, Lemma 8 says that h does not succeed on all 2-by-2-by-2 games. An example is given in its proof of a game with a PNE where h fails to converge.

Lemma 8. *No historyless uncoupled stationary strategy mapping succeeds on $\mathcal{U}(2, 2, 2)$.*

Proof of Theorem 5. Let $A = A_1 \times \dots \times A_n$. By Observation 2, it suffices to show that h does not succeed on $\mathcal{U}(|A_1|, \dots, |A_n|)$. Assume that h does succeed on $\mathcal{U}(|A_1|, \dots, |A_n|)$. If $n = 2$, $|A_1|, |A_2| > 2$, and h succeeds on $\mathcal{U}(k_1, k_2)$, then by repeatedly applying Lemma 6, h succeeds on $\mathcal{U}(3, 3)$. This contradicts Theorem 2. Now suppose that $n \geq 3$. If h succeeds on $\mathcal{U}(|A_1|, \dots, |A_n|)$, then by repeatedly applying Lemma 7, h succeeds on $\mathcal{U}(|A_1|, |A_2|, |A_3|)$. So by repeatedly applying Lemma 6, h succeeds on $\mathcal{U}(2, 2, 2)$. This contradicts Lemma 8.

3.2 Stochastic dynamics for $\mathcal{G}(A)$

We now turn to generic games and to describing the class of profile spaces A for which there exist historyless uncoupled strategy mappings that succeed on $\mathcal{G}(A)$. Theorem 3 tells us that h succeeds on two-player generic games. In fact, h also succeeds on three-player generic games where one player has only two options.

Theorem 9. *Let A be a three-player profile space such that one player has only two actions. Then h succeeds on all generic games (A, U) .*

The proof of this theorem relies partially on an analogy between a k -by- ℓ -by-2 generic game and a $k\ell$ -by-2 game that might not be generic. This requires the following technical lemma showing that under h , two players in a generic game sometimes behave similarly to a single player.

Lemma 10. *Let $k, \ell \in \mathbb{N}$, and let $U \in \mathcal{G}(k, \ell)$ be a game in which neither player has a strictly dominant action. For every $a, b \in A$ such that a is not a PNE for U , b is h^U -reachable from a .*

Proof of Theorem 9. Let $A = \{1, \dots, k\} \times \{1, \dots, \ell\} \times \{0, 1\}$ for some $\ell, k \in \mathbb{N}$. Let $U \in \mathcal{G}(A)$ and $a = (a_1, a_2, a_3) \in A$. All PNE are absorbing states under h , so it will suffice to show there is some PNE that is h^U -reachable from a .

Let $A' = \{1, \dots, k\} \times \{1, \dots, \ell\}$, and consider the games $U^0 = (u_1^0, u_2^0)$ and $U^1 = (u_1^1, u_2^1) \in \mathcal{G}(A')$ defined by

$$\begin{aligned} u_i^0(x_1, x_2) &= u_i(x_1, x_2, 0) \\ u_i^1(x_1, x_2) &= u_i(x_1, x_2, 1) \end{aligned}$$

for every $x_1 \in \{1, \dots, k\}$, $x_2 \in \{1, \dots, \ell\}$, and $i \in \{0, 1\}$. In this proof we will repeatedly use the fact that over any finite number of steps, w.p.p. player 3 doesn't move, so if $(y_1, y_2) \in A'$ is h^{U^0} -reachable from $(x_1, x_2) \in A'$, then $(y_1, y_2, 0) \in A$ is h^U -reachable from $(x_1, x_2, 0) \in A$, and similarly for h^{U^1} .

Claim. If either player has a strictly dominant action in U^0 or U^1 , then some PNE is h^U -reachable from a .

Thus we may assume that neither player has a strictly dominant action in U^0 or in U^1 . Consider a two-player game $\hat{U} = (\hat{u}_1, \hat{u}_2)$ on $\hat{A} = (\{1, \dots, k\} \times \{1, \dots, \ell\}) \times \{0, 1\}$ given by

$$\begin{aligned} \hat{u}_1(x) &= \begin{cases} 1 & \text{if } (x_1, x_2) \text{ is a PNE for } U^{x_3} \\ 0 & \text{otherwise} \end{cases} \\ \hat{u}_2(x) &= u_3((x_1, x_2, x_3)), \end{aligned}$$

for every $x = ((x_1, x_2), x_3) \in \hat{A}$. Note that unlike U , this game is not necessarily generic. By Theorem 4, some PNE $\hat{p} = ((p_1, p_2), p_3)$ for \hat{U} is $h^{\hat{U}}$ -reachable from $\hat{a} = ((a_1, a_2), a_3)$.

Now let $\hat{x} = ((x_1, x_2), x_3)$ and $\hat{y} = ((y_1, y_2), y_3) \in \hat{A}$ such that w.p.p. $\hat{y} = h^{\hat{U}}(\hat{x})$. If $x_3 \neq y_3$, then $x_3 \notin BR_2^{\hat{U}}(\hat{x})$, so $x_3 \neq BR_3^U(x)$. Thus w.p.p. $h^U(x) = (x_1, x_2, y_3)$. Since $BR_3^U(x) \neq x_3 \neq y_3$ and $|A_3| = 2$, we must have $BR_3^U(x) = y_3$, so if (x_1, x_2) is a PNE for U^{y_3} , then (x_1, x_2, y_3) is a PNE for U . Otherwise, by Lemma 10 (y_1, y_2) is $h^{U^{x_3}}$ -reachable from (x_1, x_2) , so $y = (y_1, y_2, y_3)$ is h^U -reachable from (x_1, x_2, y_3) and therefore from x .

Applying this to the each step on the path by which \hat{p} is $h^{\hat{U}}$ -reachable from \hat{a} , we see that either $p = (p_1, p_2, p_3)$ (which is a PNE for U) is h^U -reachable from a , or some other PNE for U is encountered in this process and thus h^U -reachable from a .

In fact, two-player and 2-by- k -by- ℓ are the only sizes of generic games on which h always succeeds.

Theorem 11. *Let A be a profile space. If A has more than three players, or if every player has more than two actions, then no historyless uncoupled stationary strategy mapping succeeds on all generic games (A, U) .*

Before proving this theorem, we present two lemmas. Lemma 12 says that h does not succeed on all 2-by-2-by- k -by- ℓ generic games. It is proved by giving an example of such a game.

Lemma 12. *For every $k, \ell \geq 2$, h does not succeed on $\mathcal{G}(2, 2, k, \ell)$.*

Lemma 13 says that h doesn't succeed on all three-player generic games in which all players have at least three actions. This is demonstrated by simple modifications of the 3-by-3-by-3 game used by Hart and Mas-Colell in their proof of Theorem 2.

Lemma 13. *For every $k_1, k_2, k_3 \geq 3$, h does not succeed on $\mathcal{G}(k_1, k_2, k_3)$*

Proof of Theorem 11. By Observation 2, it suffices to show that h does not succeed on $\mathcal{G}(|A_1|, \dots, |A_n|)$. Assume for contradiction that h does succeed on $\mathcal{G}(|A_1|, \dots, |A_n|)$. If $n = 3$ and h succeeds on $\mathcal{G}(|A_1|, |A_2|, |A_3|)$, then by Lemma 13 we cannot have $|A_1|, |A_2|, |A_3| > 2$. If $n = 4$ and h succeeds on $\mathcal{G}(|A_1|, \dots, |A_4|)$, then by Lemma 12 there are distinct $i, j, k \in \{1, 2, 3, 4\}$ such that $|A_i|, |A_j|, |A_k| > 2$. But by Lemma 7, h succeeds on $\mathcal{G}(|A_i|, |A_j|, |A_k|)$, contradicting lemma 13. If $n > 4$ and h succeeds on $\mathcal{G}(|A_1|, \dots, |A_n|)$, then by repeatedly applying Lemma 12, h succeeds on $\mathcal{G}(|A_1|, \dots, |A_4|)$, which we have already shown to be impossible.

4 Deterministic uncoupled dynamics

Both h and the strategy mapping used by Hart and Mas-Colell [8] to prove Theorem 1 are variations on random search. For deterministic dynamics, an exhaustive search requires more structure, and the challenge for deterministic players in short-recall uncoupled dynamics is in keeping track of their progress in the search.

We show that there are successful 3-recall deterministic dynamics by using repeated profiles to coordinate.

Theorem 14. *For every profile space A , there exists a deterministic uncoupled 3-recall stationary strategy mapping that succeeds on all games (A, U) .*

Proof. Let $n \geq 2$, $k_1, \dots, k_n \geq 2$, and $A = \{1, \dots, k_1\} \times \dots \times \{1, \dots, k_n\}$. It suffices to show that such a strategy mapping exists for $\mathcal{U}(A)$. Let $\sigma : A \rightarrow A$ be a cyclic permutation on the profiles. We write $\sigma_i(a)$ for the action of player i in $\sigma(a)$. Let $f : \mathcal{U}(A) \rightarrow \mathcal{F}(A)$ be the strategy mapping such that, for every game $U \in \mathcal{U}(A)$, player $i \in \{1, \dots, n\}$, and state $x = (a, b, c) \in A^3$,

$$f_i^U(x) = \begin{cases} c_i & \text{if } b = c \text{ and } c_i \in BR_i(c) \\ \min BR_i(c) & \text{if } b = c \text{ and } c_i \notin BR_i(c) \\ \sigma_i(a) & \text{if } a = b \neq c \\ c_i & \text{otherwise.} \end{cases}$$

Informally, the players use repetition to keep track of which profile is the current “PNE candidate” in each step. If a profile has just been repeated, then it is the current candidate, and each player plays a best reply to it, with a preference against moving. If the players look back and see that some profile a was repeated in the past but then followed by a different profile, they infer that a was rejected as a candidate and move on by playing a ’s successor, $\sigma(a)$. Otherwise the players repeat the most recent profile, establishing it as the new candidate. We call these three types of states *query*, *move-on*, and *repeat* states, respectively. Here “query” refers to asking each player for one of its best replies to b .

Let $U \in \mathcal{U}(A)$ be a game with at least one PNE. We wish to show that f^U guarantees convergence to a PNE. Let $x = (a, b, c) \in A^3$, and let y be the next state $(b, c, f^U(x))$. If x is a repeat state, then $y = (b, c, c)$, which is a query state. If x is a move-on state, then $b \neq c$, and $y = (b, c, \sigma(a))$. If $c = \sigma(a)$, then this is a query state; otherwise, it’s a repeat state, which will be followed by the query state $(c, \sigma(a), \sigma(a))$. Thus every non-query state will be followed within two steps by a query state.

Now let $x = (a, b, b) \in A^3$ be a query state, and let y and z be the next two states. If b is a PNE, then $y = (b, b, b)$, which is an absorbing state. Otherwise, $y = (b, b, c)$ for some $c \neq b$, so y is a move-on state, which will be followed by a query state $(b, \sigma(b), \sigma(b))$ or $(c, \sigma(b), \sigma(b))$ within two steps. Let p be a PNE for U . Since σ is cyclic, $p = \sigma^r(b)$ for some $r \in \mathbb{N}$. So (p, p, p) is reachable from x unless $\sigma^s(b)$ is a PNE for some $s < r$. It follows that f^U guarantees convergence to a PNE, so f succeeds on $\mathcal{U}(A)$. \square

Recall that Lemma 6 says that in the stochastic setting, adding actions to a profile space A does not make success on $\mathcal{U}(A)$ any easier. In light of that result, it is perhaps surprising that we can improve on the above bound when every player has sufficiently many actions.

Theorem 15. *If A is a profile space in which every player has at least four actions, then there exists a 2-recall deterministic uncoupled stationary strategy mapping that succeeds on all games (A, U) .*

Proof. Let $n \geq 2$, $k_1, \dots, k_n \geq 4$, and $A = \{1, \dots, k_1\} \times \dots \times \{1, \dots, k_n\}$. It suffices to show that such a strategy mapping exists for $\mathcal{U}(A)$.

Define a permutation $\sigma : A \rightarrow A$ such that for every $a \in A$, $\sigma(a)$ is a ’s lexicographic successor. Formally, $\sigma(a) = (\sigma_1(a), \dots, \sigma_n(a))$ where for $i = 1, \dots, n-1$,

$$\sigma_i(a) = \begin{cases} a_i + 1 \bmod k_i & \text{if } a_j = k_j \text{ for every } j \in \{i+1, \dots, n\} \\ a_i & \text{otherwise,} \end{cases}$$

and $\sigma_n(a) = a_n + 1 \bmod k_n$. Observe then that σ is cyclic, and for each player i and $a \in A$, we have

$$\sigma_i(a) - a_i \bmod k_i \in \{0, 1\}.$$

We now describe a strategy mapping $f : \mathcal{U}(A) \rightarrow \mathcal{F}(A)$. To each $U \in \mathcal{U}$, f assigns the strategy vector f^U defined as follows. At state $x = (a, b) \in A^2$, f^U differentiates between three types of states, each named according to the event it prompts:

- *move-on*: If $a \neq b$ and $a_j - b_j \bmod k_j \in \{0, 1\}$ for every $j \in \{1, \dots, n\}$, then the players “move on” from a , in the sense that each player i plays $\sigma_i(a)$, giving $f^U(x) = \sigma(a)$.
- *query*: If $b_j - a_j \bmod k_j \in \{0, 1, 2\}$, then we “query” each player’s utility function to check whether it is U -best-replying at b . Each player i answers by playing b_i if it is best-replying and $b_i - 1 \bmod k_i$ if it is not. So at query states,

$$f_i^U(x) = \begin{cases} b_i & \text{if } b_i \in BR_i(b) \\ b_i - 1 \bmod k_i & \text{otherwise,} \end{cases}$$

for $i = 1, \dots, n$.

- *repeat*: Otherwise, each player i “repeats” by playing b_i , giving $f^U(x) = b$.

Notice that because $k_1, \dots, k_n \geq 4$, it is never the case that both $a_j - b_j \bmod k_j \in \{0, 1\}$ and $b_j - a_j \bmod k_j \in \{0, 1, 2\}$. Thus the conditions for the *move-on* and *query* types are mutually exclusive, and the three state types are all disjoint.

The state following $x = (a, b)$ is $y = (b, f^U(x))$. If x is a move-on state, then $y = (b, \sigma(a))$. Since for every player i , $a_i - b_i \bmod k_i \in \{0, 1\}$ and $\sigma(a)_i - a_i \bmod k_i \in \{0, 1\}$, we have $\sigma_i(a) - b_i \bmod k_i \in \{0, 1, 2\}$, so y is a query state. If x is instead a query state, then $b_i - f_i^U(x) \bmod k_i \in \{0, 1\}$ for every player i , so y is a move-on state unless $b = f^U(x)$, in which case $y = (b, b)$ is a query state. But if $b = f^U(x)$ and x was a query state, then $b_i \in BR_i(b)$ for every player i , i.e., b is a PNE. Finally, if x is a repeat state, then $y = (b, b)$ is a query state.

Thus move-on states and repeat states are always followed by query states, and ask-all states are never followed by repeat states. We conclude that with the possible exception of the initial state, every state will be a move-on or query state, and no two consecutive states will be move-on states. In particular, some query state is reachable from every initial state.

For any query state $x = (a, b)$, x will be followed by (b, b) if and only if b is a PNE, and (b, b) is an absorbing state for every PNE b . If b is not a PNE, then x will be followed will be a move-on state (b, c) , for some $c \in A$. This will be followed by the *query* state $(c, \sigma(b))$. Continuing inductively, since σ is cyclic, unless the players converge to a PNE, they will examine every profile $v \in A$ with a query state of the form (u, v) . Thus for every game U with at least one PNE, f^U guarantees convergence to a PNE, i.e., f succeeds on $\mathcal{U}(A)$. \square

While there are deterministic uncoupled 2-recall dynamics that succeed on at least some classes that require 2-recall in the stochastic setting, historyless dynamics of this type fail on $\mathcal{U}(A)$ for every profile space A .

Theorem 16. *For every profile space A , no deterministic uncoupled historyless stationary strategy mapping succeeds on all games (A, U) .*

5 Future Directions

It remains open to determine tight bounds on the minimum recall of successful deterministic uncoupled dynamics for every profile space, analogous to those

given in Section 3 for stochastic dynamics. In particular, we make the following conjecture.

Conjecture 1. *There exists a profile space A such that no deterministic uncoupled 2-recall strategy mapping succeeds on all games (A, U) .*

The same questions answered in this work may naturally be asked for other important classes of games (e.g., symmetric games) and other equilibrium concepts, especially mixed Nash equilibrium. More generally, the resources (e.g., recall, memory) required by uncoupled self-stabilizing dynamics in asynchronous environments should be investigated.

References

1. Y. Babichenko. Uncoupled automata and pure Nash equilibria. *Int. J. Game Theory*, 39(3):483–502, 2010.
2. Y. Babichenko. Completely uncoupled dynamics and Nash equilibria. *Games and Economic Behavior*, 76(1):1–14, 2012.
3. S. Dolev. *Self-Stabilization*. MIT Press, 2000.
4. D. Fudenberg and J. Tirole. *Game theory*. MIT Press, 1991.
5. J. Y. Halpern. Computer science and game theory: A brief survey. <http://arxiv.org/abs/cs/0703148>, 2007.
6. S. Hart and Y. Mansour. How long to equilibrium? The communication complexity of uncoupled equilibrium procedures. *Games and Economic Behavior*, 69(1):107–126, 2010.
7. S. Hart and A. Mas-Colell. Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93(5):1830–1836, 2003.
8. S. Hart and A. Mas-Colell. Stochastic uncoupled dynamics and Nash equilibrium. *Games and Economic Behavior*, 57(2):286–303, 2006.
9. S. Hart and A. Mas-Colell. *Simple Adaptive Strategies: From Regret-matching to Uncoupled Dynamics*. World Scientific Publishing Co., Inc., River Edge, NJ, USA, 2012.
10. A. D. Jaggard, M. Schapira, and R. N. Wright. Distributed computing with adaptive heuristics. In B. Chazelle, editor, *ICS*, pages 417–443. Tsinghua University Press, 2011.
11. N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
12. B. S. R. Pradelski and H. P. Young. Learning efficient Nash equilibria in distributed systems. *Games and Economic Behavior*, 75(2):882–897, 2012.
13. H. P. Young. Learning by trial and error. *Games and Economic Behavior*, 65(2):626–643, 2009.