Self-stabilizing uncoupled dynamics*

Aaron D. Jaggard[1], Neil Lutz[2], Michael Schapira[3], and Rebecca N. Wright[4]

[1] Formal Methods Section (Code 5543), U.S. Naval Research Laboratory, Washington, DC 20375, USA. aaron.jaggard@nrl.navy.mil
[2] Department of Computer Science, Rutgers University, Piscataway, NJ 08854, USA. njlutz@cs.rutgers.edu
[3] School of Computer Science and Engineering, Hebrew University of Jerusalem, Jerusalem 91904, Israel. schapiram@huji.ac.il
[4] Department of Computer Science, 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. For uncoupled players, recognizing an equilibrium is a distributed computational task. Self-stabilizing uncoupled dynamics, then, capture both resilience to failure and distribution of knowledge. We initiate the study of these dynamics by analyzing their strength in a bounded-recall synchronous setting. We determine, for every size of game, the minimum number of periods of play that stochastic agents must recall in self-stabilizing uncoupled dynamics, and also in the special case when the game is guaranteed to have unique best replies. For deterministic agents, we give 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. We also prove that for uncoupled deterministic agents a single step of recall is insufficient to achieve self-stabilization, regardless of game size.

1 Introduction

Self-stabilization is a failure-resilience property that is central to distributed computing theory and is the subject of extensive research [3]. 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 [4,10]. In particular, game theory researchers have investigated dynamics that are uncoupled, in the sense that the way each player plays the game is independent of the other players' payoffs [8]. Here we bring together these two research areas and initiate the study of *self-stabilizing uncoupled dynamics*.

In this work we focus our investigation on a bounded-recall synchronous setting. We consider self-stabilization in a multi-agent distributed system in

* Research supported in part by National Science Foundation Grant 1101690.

which, at each timestep, the agents act as 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. The dynamics we study will have *bounded recall*, so the state of this system at any time consists of the r most recent action profiles, for some finite r. A stable state is one where the same action profile appears in r consecutive timesteps and is a *pure Nash equilibrium* (PNE), i.e., no player could benefit from unilaterally choosing a different action. These game dynamics self-stabilize if from every state the players are guaranteed to converge to a PNE.

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. Unlike these *behavioral* restrictions on the players, uncoupledness is an *informational* restriction: 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. The concept of uncoupled game dynamics was introduced by Hart and Mas-Colell [6].

If uncoupledness is the only restriction on the dynamics, then the players can find a PNE through a straightforward exhaustive search, but Hart and Mas-Colell [6] demonstrated in a continuous-time setting that deterministic uncoupled dynamics fail to reach a stable state for some games that have PNE if we additionally require that the dynamics are *historyless*, in the sense that the state space of the system is identical to the action profile space of the game. This suggests the central question of the present work:

> On a given class of games, how much recall do uncoupled players need in order to self-stabilize whenever a PNE exists?

Throughout this paper, we say that dynamics *succeed* on a class of games if they self-stabilize whenever a PNE exists.

This question was answered in part by Hart and Mas-Colell [7] when they proved in a discrete-time setting that even when players are allowed randomness, historyless uncoupled dynamics cannot 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" action), historyless uncoupled dynamics do not succeed on games played by three three-action players. They also gave positive results, proving that historyless uncoupled dynamics *can* 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 there are successful stochastic uncoupled dynamics for all games over every action profile space.

Our results We show in Section 3 that historyless uncoupled dynamics can also succeed on all two-player games with a two-action player and on all three-player generic games with a two-action player (Theorems 6 and 11). In both cases, we prove that these results do not hold for any larger size of game (Theorems 7 and 13). Combined with the results of Hart and Mas-Colell [7], this tells us, for any action profile space, the exact minimum recall needed 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 give 3-recall deterministic uncoupled dynamics that succeed on every game (Theorem 16) and 2-recall deterministic uncoupled dynamics that succeed on every game in which every player has at least four actions (Theorem 17). We also prove that historyless deterministic uncoupled dynamics cannot succeed on all games over *any* action profile space (Theorem 18).

Related work In addition to the results mentioned above, Hart and Mas-Colell also addressed convergence to mixed Nash equilibria by bounded-recall uncoupled dynamics [7]. 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]. Hart and Mansour [5] analyzed the time to convergence for uncoupled dynamics. Jaggard, Schapira, and Wright [9] investigated convergence to pure Nash equilibria by game dynamics that are distributed in the sense of being asynchronous, rather than uncoupled.

2 Definitions

Games Let $n \in \mathbb{N}$ and $(k_1, ..., k_n) \in \mathbb{N}^n$, with each $k_i \geq 2$. A *game* of *size* $(k_1, ..., k_n)$ is a pair (A, U), where $A = A_1 \times ... \times A_n$ such that each $|A_i| = k_i$, and $U = (u_1, ..., u_n)$ is an n-tuple of functions $u_i : A \to \mathbb{R}$. A_i and u_i are the *action set* and *utility function* of *player* i. 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 the class of all U such that each u_i takes A_i as input, so $(A, \mathcal{U}(A))$ 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 $i \in \{1, ..., n\}$ and $a = (a_1, ..., a_n) \in A$, we say that player i is U-*best-replying* at a if $u_i(a) \geq u_i((a_1, ..., a_i', ..., a_n))$ for every $a_i' \in A_i$. We define the set of U-*best-replies* for player i at a,

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

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 $i \in \{1, ..., n\}$ 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 $(A, U) \in (A, \mathcal{U}(A))$ is *generic* if every player's best-replies are unique, i.e., if for every $a \in A$ and $i \in \{1, ..., n\}$, $|BR_i^U(a)|$=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. $(A, \mathcal{G}(A))$ 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 $a^{(t)} = (a_1^{(t)}, ..., a_n^{(t)})$. The *stage game* $(A, U) \in (A, \mathcal{U}(A))$ is then played: each player i simultaneously selects a new action $a_i^{(t+1)}$ by applying an *r-recall stationary strategy* $f_i^U : A^r \to A_i$, where $r \in \mathbb{N}$ and A^r is the cartesian product of A with itself r times. The strategy f_i^U, which is *stationary* in the sense that it does not depend on t, will take as input $(a^{(t-r+1)}, ..., a^{(t)})$, the r most recently observed profiles. We call this r-tuple the *state* at time t. Unless we specify that it is *deterministic*, f_i^U is permitted to be a random variable. The terms 1-recall and *historyless* are interchangeable. A *strategy vector* is an n-tuple $f^U = (f_1^U, ..., f_n^U)$, 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 $f : \mathcal{U}(A) \to \mathcal{F}(A)$ 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 $f_1, ..., f_n$ where each f_i maps utility functions on A to strategies for A, such that $f_i(u_i) \equiv f_i^U$ for $i = 1, ..., n$. If f_i^U is stationary, deterministic, or r-recall for $i = 1, ..., n$, then f^U is also. If every f^U has any of those properties, then f does also.

Now let $x = (x^{(1)}, ..., x^{(r)}) \in A^r$, and let f^U be an r-recall strategy vector. For $T \geq r$, a *partial f^U-run* for T steps is a sequence of profiles $\{a^{(t)}\}_{t=0}^T \subseteq A$ that results with positive probability (w.p.p.) from beginning at state $x = (a^{(1)}, ..., a^{(r)})$ and repeatedly applying f^U. That is, for every $r < t \leq T$,

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

An *f^U-run* is a sequence of profiles $\{a^{(t)}\}_{t=0}^\infty$ 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 $T \geq r$ and a partial f^U-run $\{a^{(t)}\}_{t=0}^T$ such that $x = (a^{(1)}, ..., a^{(r)})$ and $y = (a^{(T+1-r)}, ..., a^{(T)})$. The state x is an *f^U-absorbing state* if for every f^U-run $\{a^{(t)}\}_{t=0}^\infty$ beginning from x, $(a^{(t+1)}, ..., a^{(t+r)}) = x$ for every $t \in \mathbb{N}$. Notice that any f^U-absorbing state $x = (a^{(1)}, ..., a^{(r)})$ must have $a^{(1)} = ... = a^{(r)}$. 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 $\{a^{(t)}\}_{t \in \mathbb{N}}$ *converges* to a profile a if there some $T \in \mathbb{N}$ such that $a^{(t)} = a$ for every $t \geq T$. Suppose that from every $x \in A^r$, some f^U-absorbing state is f^U-reachable. Then almost every run of f^U converges. We say that f *self-stabilizes* on game (A, U) if almost every f^U-run converges to a PNE. 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 $(A, U) \in (A, \mathcal{C}(A))$, then f *succeeds* on $\mathcal{C}(A)$. Where f does not succeed it *fails*.

Let $A = A_1 \times ... \times A_n$ and $B = B_1 \times ... \times B_n$ be profile spaces of the same size, in the sense that there is some permutation π on $\{1, ..., n\}$ such that $(|A_1|, ..., |A_n|) = (|B_{\pi(1)}|, ..., |B_{\pi(n)}|)$. 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

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

and we say that f succeeds on $\mathcal{C}(|A_1|, ..., |A_n|)$ 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 can succeed on all games $(A, U) \in (A, \mathcal{U}(A))$ or all generic games $(A, U) \in (A, \mathcal{G}(A))$. Hart and Mas-Colell [7] 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 [7]) *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 [7]) *There is no uncoupled historyless stationary strategy mapping that succeeds on all three-player games, or on all 3-by-3-by-3 generic games.*

Theorem 3 (Hart and Mas-Colell [7]) *For any two-player profile space A, there is an uncoupled historyless stationary strategy mapping that succeeds on all 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 $h : \mathcal{U}(A) \to \mathcal{F}(A)$, defined as follows. Let $U = (u_1, ..., u_n) \in \mathcal{U}(A)$. Then $h(U) = (h_1^U, ..., h_n^U)$, where for $i \in \{1, ..., n\}$, $h_i^U : A \to A_i$ is given by

$$\Pr\left(h^U(a_i) = a_i \,|\, a_i \in BR_i(a)\right) = 1$$
$$\Pr\left(h^U(a_i) = b_i \,|\, a_i \notin BR_i(a)\right) = 1/n,$$

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.

Let A be a profile space, and let f be an uncoupled historyless stationary strategy mapping for A. In their proof of Theorem 2, Hart and Mas-Colell make the following observation.

Observation 4 (Hart and Mas-Colell [7]) *If f succeeds on all generic games (A, U), then two conditions hold for every game (A, U) and $a = (a_1, ..., a_n) \in A$. 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 5 (Hart and Mas-Colell [7]) *If any historyless uncoupled 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 uncoupled historyless strategy mappings can succeed on every game, or equivalently (by Observation 5), the A for which h succeeds on $\mathcal{U}(A)$. The proof that h succeeds on 2-by-k games proceeds by simple case checking but may be a useful warmup for working with these dynamics.

Theorem 6 *For every two-player profile space A in which one player has only two actions, h succeeds on all games (A, U).*

Proof. Let $k \geq 2$. It suffices to show that h succeeds on $\mathcal{U}(2, k)$. Let $A = \{1, 2\} \times \{1, ..., k\}$ and $U = (u_1, u_2) \in \mathcal{U}(A)$. Suppose that U has at least one PNE, and recall that every PNE for U is an h^U-absorbing state. Let $a = (a_1, a_2) \in A$, and consider four cases.

1. Player 1 is best-replying at a and $a_1 = p_1$ for some PNE $p = (p_1, p_2)$. Then either player 2 is also best-replying and a is a PNE, or $h_2^U(a) = p_2$ w.p.p., so $h^U(a)$ is a PNE w.p.p.
2. Player 1 is not best-replying at a and there is no PNE p such that $a_1 = p_1$. Then w.p.p. $h_1^U(a) \neq a_1$ and $h_2^U(a) = a_2$. Since we assumed that U has a PNE, $h^U(a)$ is then an instance of case 1.
3. Player 1 is best-replying at a and there is no PNE p such that $a_1 = p_1$. If player 2 is also best-replying, then a is a PNE. Otherwise, w.p.p. $h_2^U(a) \in BR_2(a)$, but $h^U(a)$ cannot be a PNE since $h_1^U(a) = a_1$. So player 1 is not best-replying at $h^U(a)$, i.e., $h^U(a)$ is an instance of case 2.

4. Player 1 is not best-replying at a and $a_1 = p_1$ for some PNE $p = (p_1, p_2)$. Then w.p.p. $h_1^U(a) \neq a_1$ and $h_2^U(a) = a_2$, so player 1 is best-replying at $h^U(a)$, meaning that it is an instance of case 1 or 3.

We conclude that from every state $a \in A$, some PNE for U is h^U-reachable from a. Thus h succeeds on $\mathcal{U}(A)$. $\square$

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

Theorem 7 *Let A be a profile space. Unless A has only two players and one of those players has only two actions, no historyless uncoupled strategy mapping succeeds on all games (A, U).*

We give three lemmas that will be used in the proof of Theorem 7. Their full proofs are in the appendix. Informally, Lemma 8 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 8 *Let $n \geq 2$, $k_1, ..., k_n \geq 2$, and $i \in \{1, ..., n\}$. If h succeeds on $\mathcal{U}(k_1, ..., k_i + 1, ..., k_n)$, then h succeeds on $\mathcal{U}(k_1, ..., k_i, ..., k_n)$.*

Lemma 9 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 9 *Let $n \geq 2$ and $k_1, ..., k_n, k_{n+1} \geq 2$. If h succeeds on $\mathcal{U}(k_1, ..., k_n, k_{n+1})$, then h succeeds on $\mathcal{U}(k_1, ..., k_i, ..., k_n)$. The same is true if we replace $\mathcal{U}$ with $\mathcal{G}$.*

Finally, Lemma 10 says that h does not succeed on all 2-by-2-by-2 games. An example is given in its proof of a game where h fails.

Lemma 10 *No historyless uncoupled strategy mapping succeeds on $\mathcal{U}(2, 2, 2)$.*

Proof of Theorem 7. Let $A = A_1 \times ... \times A_n$. By Observation 5, it suffices to show that h does not succeed on $\mathcal{U}(|A_1|, ..., |A_n|)$. Assume that h does succeed on $\mathcal{U}(|A_1|, ..., |A_n|)$. If $n = 2$, $|A_1|, |A_2| > 2$, and h succeeds on $\mathcal{U}(k_1, k_2)$, then by repeatedly applying Lemma 8, h succeeds on $\mathcal{U}(3, 3)$. This contradicts Theorem 2. Now suppose that $n \geq 3$. If h succeeds on $\mathcal{U}(|A_1|, ..., |A_n|)$, then by repeatedly applying Lemma 9, h succeeds on $\mathcal{U}(|A_1|, |A_2|, |A_3|)$. So by repeatedly applying Lemma 8, h succeeds on $\mathcal{U}(2, 2, 2)$. This contradicts Lemma 10. $\square$

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 h (or any historyless uncoupled strategy mapping) can succeed on $\mathcal{G}(A)$. Theorem 3 tells us that h succeeds on two-player generic games. That also succeeds on three-player generic games where one player has only two options.

Theorem 11 *Let A be a three-player profile space such that one player has only three 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 kl-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 12 *Let $k, l \in \mathbb{N}$, and let $U \in \mathcal{G}(k, l)$ 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 11. Let $A = \{1, ..., k\} \times \{1, ..., l\} \times \{0, 1\}$ for some $l, k \in \mathbb{N}$. 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 $A' = \{1, ..., k\} \times \{1, ..., l\}$, 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

$$u_i^0(x_1, x_2) = u_i(x_1, x_2, 0)$$
$$u_i^1(x_1, x_2) = u_i(x_1, x_2, 1)$$

for every $x_1 \in \{1, ..., k\}$, $x_2 \in \{1, ..., l\}$, 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.

This claim is proved in the appendix. Thus we may assume that neither player has a strictly dominant action in U^0 or in U^1. Consider a two-player game $\widehat{U} = (\widehat{u}_1, \widehat{u}_2)$ on $\widehat{A} = (\{1, ..., k\} \times \{1, ..., l\}) \times \{0, 1\}$ given by

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

for every $x = ((x_1, x_2), x_3) \in \widehat{A}$. Note that unlike U, this game is not necessarily generic. By Theorem 6, some PNE $\widehat{p} = ((p_1, p_2), p_3)$ for $\widehat{U}$ is $h^{\widehat{U}}$-reachable from $\widehat{a} = ((a_1, a_2), a_3)$.

Now let $\widehat{x} = ((x_1, x_2), x_3)$ and $\widehat{y} = ((y_1, y_2), y_3) \in \widehat{A}$ such that w.p.p. $\widehat{y} = h^{\widehat{U}}(\widehat{x})$. If $x_3 \neq y_3$, then $x_3 \notin BR_2^{\widehat{U}}(\widehat{x})$, 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 12 (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 $\widehat{p}$ is $h^{\widehat{U}}$-reachable from $\widehat{a}$, 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 13 *Let A be a profile space. If A has more than three players, or if every player has more than two actions, then no uncoupled historyless stationary strategy mapping can succeed on all generic games (A, U).*

Before proving this theorem, we present two lemmas whose full proofs are in the appendix. Lemma 14 says that h fails on some 2-by-2-by-k-by-l generic games. It is proved by giving an example of such a game.

Lemma 14 *For every $k, l \geq 2$, h does not succeed on $\mathcal{G}(2, 2, k, l)$.*

Lemma 15 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 15 *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 13. By Observation 5, if suffices to show that h does not succeed on $\mathcal{G}(|A_1|, ..., |A_n|)$. Assume for contradiction that h does succeed on $\mathcal{G}(|A_1|, ..., |A_n|)$. If $n = 3$ and h succeeds on $\mathcal{G}(|A_1|, |A_2|, |A_3|)$, then by Lemma 15 we cannot have $|A_1|, |A_2|, |A_3| > 2$. If $n = 4$ and h succeeds on $\mathcal{G}(|A_1|, ..., |A_4|)$, then by Lemma 14 there are distinct $i, j, k \in \{1, 2, 3, 4\}$ such that $|A_i|, |A_j|, |A_k| > 2$. But by Lemma 9, h succeeds on $\mathcal{G}(|A_i|, |A_j|, |A_k|)$, contradicting lemma 15. If $n > 4$ and h succeeds on $\mathcal{G}(|A_1|, ..., |A_n|)$, then by repeatedly applying Lemma 14, h succeeds on $\mathcal{G}(|A_1|, ..., |A_4|)$, which we have already shown to be impossible. □

4 Deterministic uncoupled dynamics

Both h and the strategy mapping used by Hart and Mas-Colell [7] 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.

4.1 Positive results

We show that the players can succeed on all games with 3-recall by using repeated profiles to coordinate.

Theorem 16 *For every profile space A, there exists a deterministic uncoupled 3-recall stationary strategy mapping that succeeds on all games (A, U).*

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

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

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 $\mathcal{U}(A)$. □

Recall that Lemma 8 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 17 *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 $n \geq 2$, $k_1, ..., k_n \geq 4$, and $A = \{1, ..., k_1\} \times ... \times \{1, ..., k_n\}$. It suffices to show that such a strategy mapping exists for $\mathcal{U}(A)$.

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

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

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 $f : \mathcal{U}(A) \to \mathcal{F}(A)$. 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 $j \in \{1, ..., n\}$, 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 $i = 1, ..., n$.
- *repeat*: Otherwise, each player i "repeats" by playing b_i, giving $f^U(x) = b$.

Notice that because $k_1, ..., k_n \geq 4$, 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 $\mathcal{U}(A)$. □

While deterministic uncoupled 2-recall dynamics can 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 18 *For every profile space A, no deterministic uncoupled historyless stationary strategy mapping succeeds on all games (A, U).*

Proof. See appendix.

5 Open problems

We would be interested to see 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. 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 may 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 (3. pr.).* MIT Press, 1991.
5. 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.
6. S. Hart and A. Mas-Colell. Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93(5):1830–1836, 2003.
7. S. Hart and A. Mas-Colell. Stochastic uncoupled dynamics and Nash equilibrium. *Games and Economic Behavior*, 57(2):286–303, 2006.
8. 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.
9. 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.
10. N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory.* Cambridge University Press, New York, NY, USA, 2007.

Appendix

Proof of Lemma 8. Let

$$A = \{1, ..., k_1\} \times ... \times \{1, ..., k_i\} \times ... \times \{1, ..., k_n\},$$
$$A' = \{1, ..., k_1\} \times ... \times \{1, ..., k_i + 1\} \times ... \times \{1, ..., k_n\}.$$

Suppose that h succeeds on $\mathcal{U}(A')$. For each $U = (u_1, ..., u_n) \in \mathcal{U}(A)$, define another game $U' = (u'_1, ..., u'_n) \in \mathcal{U}(A')$ such that for every $j \in \{1, ..., n\}$ and $a \in A$,

$$u'_j(a) = u_j(a),$$
$$u'_j(a_1, ..., k_i + 1, ..., a_n) = u_j(a_1, ..., k_i, ..., a_n).$$

Thus in U' every player is always indifferent to whether player i plays k_i or $k_i + 1$.

We now define a strategy mapping f for games on A. For every $U \in \mathcal{U}(A)$, f^U is given by

$$\Pr\left(f^U_j(a) = h^{U'}_j(a_1, ..., k_i + 1, ..., a_n) \,\middle|\, a_i = k_i\right) = 1/2,$$
$$\Pr\left(f^U_j(a) = h^{U'}_j(a) \,\middle|\, a_i = k_i\right) = 1/2, \text{ and}$$
$$\Pr\left(f^U_j(a) = h^{U'}_j(a) \,\middle|\, a_i \neq k_i\right) = 1,$$

for every $a \in A$ and $j \neq i$. That is, whenever the players see that player i has played k_i, each chooses independently at random to interpret that action either as k_i or k_i+1, then plays the action prescribed by $h^{U'}$. Player i behaves similarly under f, but we have to ensure that it's never instructed to play $k_i + 1$:

$$\Pr\left(f^U_j(a) = \min\{k_i, h^{U'}_i(a_1, ..., k_i + 1, ..., a_n)\} \,\middle|\, a_i = k_i\right) = 1/2$$
$$\Pr\left(f^U_j(a) = \min\{k_i, h^{U'}_i(a) \,\middle|\, a_i = k_i\right) = 1/2$$
$$\Pr\left(f^U_j(a) = h^{U'}_i(a) \,\middle|\, a_i \neq k_i\right) = 1.$$

Now fix $U = (u_1, ..., u_n) \in \mathcal{U}(A)$, and assume that U has at least one PNE $p \in A$. To see that p is an absorbing state for f^U, we consider two cases. First, suppose that $p_i \neq k_i$. Then p is also a PNE for U', hence p is an absorbing state for $h^{U'}$. So for $j \neq i$, $f^U_j(p) = h^{U'}_j(p) = p_j$, and $f^U_i(p) = \min\{k_i, h^{U'}_i(p)\} = \min\{k_i, p_i\} = p_i$. Now suppose instead that $p_i = k_i$. Then both p and $p' = (p_1, ..., k_i + 1, ..., p_n)$ are stable states for $h(U')$. So for $j \neq i$,

$$\begin{aligned} f^U_j(p) &\in \{h^{U'}_j(p_1, ..., k_i + 1, ..., p_n), f^{U'}_j(p)\} \\ &= \{h^{U'}_j(p'), h^{U'}_j(p)\} \\ &= \{p_j\}, \end{aligned}$$

and

$$\begin{aligned}f_i^U(p) &\in \left\{ \min\{k_i, h_i^{U'}(p_1, ..., k_i+1, ..., p_n)\}, \min\{k_i, h_i^{U'}(p)\} \right\} \\ &\subseteq \left\{ \min\{k_i, h_i^{U'}(p')\}, k_i, h_i^{U'}(p) \right\} \\ &= \left\{ \min\{k_i, p'_i\}, k_i, p_i \right\} \\ &= \{p_i\}.\end{aligned}$$

Thus p is an absorbing state for f^U.

It remains to show that f^U always reaches a PNE. Let $a \in A \subseteq A'$. Since U' has a PNE and h succeeds on $\mathcal{U}(A')$, U' has some PNE $q = (q_1, ..., q_n) \in A'$ such that q is $h^{U'}$-reachable from a. So for some $T \in \mathbb{N}$, theres is a partial $h^{U'}$-run $a^{(0)}, ..., a^{(T)}$ such that $a^{(0)} = a$ and $a^{(T)} = q$. Since q is a PNE for U', $q' = (q_1, ..., \min\{q_i, k_i\}, ..., q_n)$ is a PNE for both U and U'.

Now let $b^{(0)}, ..., b^{(T)}$ be a partial f^U-run such that $b^{(0)} = a$. Suppose, for some $0 \leq t < T$, that

$$b^{(t)} = (a_1^{(t)}, ..., \min\{a_i^{(t)}, k_i\}, ..., a_n^{(t)}).$$

Then for $j \neq i$, $b_j^{(t+1)} = f_j^U(b^{(t)}) = h_j^{U'}(a^{(t)})$ with probability at least $\frac{1}{2}$, and $b_i^{(t+1)} = f_i^U(b^{(t)}) = \min\{h_j^{U'}(a^{(t)}), k_i\}$ with probability at least $\frac{1}{2}$. So with positive probability,

$$b^{(t+1)} = (a_1^{(t+1)}, ..., \min\{a_i^{(t+1)}, k_i\}, ..., a_n^{(t+1)}).$$

By induction, $\Pr[b^{(T)} = q'] > 0$, i.e., q' is f^U-reachable from a. We conclude that f succeeds on $\mathcal{U}(A)$, and by Observation 5 it follows that h succeeds on $\mathcal{U}(A)$. □

Proof of Lemma 9. Let

$$\begin{aligned}A &= \{1, ..., k_1\} \times ... \times \{1, ..., k_n\}, \\ A' &= \{1, ..., k_1\} \times ... \times \{1, ..., k_n\} \times \{1, ..., k_{n+1}\}.\end{aligned}$$

Suppose that h succeeds on $\mathcal{U}(A')$, and for each $U = (u_1, ..., u_n) \in \mathcal{U}(A)$, define a game $U' = (u'_1, ..., u'_n, u'_{n+1}) \in \mathcal{U}(A')$ such that for every $x = (x_1, ..., x_n, x_{n+1}) \in A'$,

$$u'_i(x) = u_i((x_1, ..., x_n))$$

for each player $i \in \{1, ..., n\}$ and

$$u'_{n+1}(x) = \begin{cases} 1 & \text{if } x_{n+1} = 1 \\ 0 & \text{otherwise.} \end{cases}$$

Informally, the first n players are apathetic about player $n+1$'s action, and player $n+1$ always prefers to play 1. Notice that $x = (x_1, ..., x_n) \in A$ is a PNE for U if and only if $(x_1, ..., x_n, 1)$ is a PNE for U'.

Given a game U, we use U' to define a strategy mapping f for games on A. For each $x = (x_1, ..., x_n) \in A$ and $i \in \{1, ..., n\}$,

$$f_i^U(x) = h_i^{U'}((x_1, ..., x_n, 1)).$$

Now fix $U = (u_1, ..., u_n) \in \mathcal{U}(A)$ and $a = (a_1, ..., a_n) \in A$, and assume that U has at least one pure Nash equilibrium. Then U' does also, so letting $a' = (a_1, ..., a_n, 1) \in A'$, some PNE $p' = (p_1, ..., p_n, 1)$ for U' is $h^{U'}$-reachable from a'. We show that $p = (p_1, ..., p_n)$, which is a PNE for U, is f^U-reachable from a.

Since p' is $h^{U'}$-reachable from a', there is a partial $h^{U'}$-run $a^{(0)}, ..., a^{(T)}$, for some $T \in \mathbb{N}$, such that $a^{(0)} = a'$ and $a^{(T)} = p'$. For each $t \in \{0, ..., T-1\}$, if $a_{n+1}^{(t)} = 1$, then player $n+1$ is best-replying at $a^{(t)}$, so $a_{n+1}^{(t+1)} = 1$. Thus player 1 is playing 1 at every state in the partial run. Now let $b^{(0)}, ..., b^{(T)}$ be a partial f^U-run such that $b^{(0)} = a$. At each step t, if $a^{(t)} = (b_1^{(t)}, ..., b_n^{(t)}, 1)$, then

$$f^U(b^{(t)}) = h^{U'}(b_1^{(t)}, ..., b_n^{(t)}, 1) = h^{U'}(b^{(t)}),$$

so w.p.p. $a^{(t+1)} = (b_1^{(t+1)}, ..., b_n^{(t+1)}, 1)$. It follows that w.p.p. $a^{(T)} = p'$, i.e., $b^{(T)} = p$. Thus p is f^U-reachable from a, so f succeeds on $\mathcal{U}(A)$. By Observation 5, then, h succeeds on $\mathcal{U}(A)$.

For the second part of the lemma, simply notice that U' is generic whenever U is, thus the above argument still holds when $\mathcal{G}(A)$ is substituted for $\mathcal{U}(A)$. □

Proof of Lemma 10. Let $A = \{1, 2, 3\}$. By Observation 5 it suffices to show that h does not succeed on $\mathcal{U}(A)$. Consider the game $U = (u_1, u_2, u_3) \in \mathcal{U}(A)$ where $u_i((x, y, z))$ is the ith coordinate of $M_x[y, z]$, for

$$M_1 = \begin{bmatrix} 1,1,1 & 1,0,1 \\ 1,0,0 & 0,1,1 \end{bmatrix} \qquad M_2 = \begin{bmatrix} 0,1,0 & 0,1,1 \\ 0,0,0 & 1,0,1 \end{bmatrix}.$$

The unique PNE of U is $p = (1, 1, 1)$. Let $a \in A$ with $a_3 = 2$. Then $h_3^U(a) = 2$, since $2 \in BR_3(a)$ for every a. It follows that under h^U, if the third player initially plays 2, then it will never play 1, so p is not h^U-reachable from, for example, $(1, 1, 2)$. Thus h does not succeed on $\mathcal{U}(A)$. □

Proof of Claim from Theorem 11. Suppose that player 1 has a strictly dominant action α in U^0, and consider five cases.

1. U has a PNE $(p_1, p_2, 0)$, and $a_3 = 0$. Then player 1 is best-replying at a only if $a_1 = \alpha = p_1$, so w.p.p. $h^U(a) = (p_1, a_2, 0)$. Player 2 is best-replying at $(p_1, a_2, 0)$ only if $a_2 = p_2$, so w.p.p. $h^U(h^U(a)) = (p_1, p_2, 0)$.
2. U has a PNE $(q_1, q_2, 1)$, $a_3 = 1$, and $BR_3^U(a) = 1$.
   – If some player has a strictly dominant action in U^1, then this is symmetric to the situation described in case 1, and q is h^U-reachable from a.
   – So assume that no player has a strictly dominant action in U^1. If a is not a PNE for U, then (a_1, a_2) is not a PNE for U^1. So by Lemma 12, (q_1, q_2) is h^{U^1}-reachable from (a_1, a_2), i.e., the PNE $(q_1, q_2, 1)$ is h^U-reachable from a.

3. U has no PNE $(p_1, p_2, 0)$, and $a_3 = 0$. As in case 1, w.p.p. $h^U(a) = (\alpha, a_2, 0)$. Let $b_2 = BR_2(\alpha, a_2, 0)$. Then w.p.p. $h^U((\alpha, a_2, 0)) = (\alpha, b_2, 0)$, and player 3 is not best-replying at $(\alpha, b_2, 0)$ since it is not a PNE for U. Thus letting $b = (\alpha, b_2, 1)$, w.p.p. $h^U((\alpha, b_2, 0)) = b$, so b is h^U-reachable from a, and b is an instance of case 2.
4. U has a PNE $(q_1, q_2, 1)$, $a_3 = 1$, and $BR_3^U(a) = 0$. Then w.p.p. $h^U(a) = (a_1, a_2, 0)$, which is an instance of case 1 or 3.
5. U has no PNE $(q_1, q_2, 1)$, and $a_3 = 1$.
   – If some player has a strictly dominant action in U^1, then w.p.p. that action will be played in $h^U(a)$, and w.p.p. the other player will play its best reply to that action in the next stage. Then the first two players are playing a PNE for U^1, so player 3 is not best-replying and may play 0 in the next round, giving an instance of case 1.
   – So assume that no player has a strictly dominant action in U^1. There is some $(b_1, b_2) \in A'$ such that $BR_3(b_1, b_2, 1) = 0$, so w.p.p. $h^U(b_1, b_2, 1) = (b_1, b_2, 0)$. If (a_1, a_2) is a PNE for U^1, then player 3 is not best replying and w.p.p. $h^U(a) = (a_1, a_2, 0)$. Otherwise by Lemma 12 (b_1, b_2) is h^{U^1}-reachable from (a_1, a_2), so $(b_1, b_2, 0)$, which is an instance of case 1, is h^U-reachable from a.

It follows that some PNE for U is h^U-reachable from every $a \in A$. By symmetry, the same holds whenever either player has a strictly dominant action in either U^0 or U^1. □

Proof of Lemma 12. If $k = l = 2$, then each player either prefers to match or to mismatch the other's action, and lemma holds by routine inspection of the four possibilities. So assume $l > 2$.

Let $a, b \in A$, where a is not a PNE for U. Notice that because U is generic, A contains exactly l states where player 1 is best-replying and k states where player 2 is best-replying, so there are at most $k + l$ states where either player is best-replying. And for any $x, y \in A$, if neither player is best-replying at x, then $h^U(x) = y$ with probability $\frac{1}{kl}$. Hence it suffices to show that more than $k + l$ distinct states in A are reachable from a.

If player 2 is best-replying at a, then since player 2 has no dominant action, player 1 has some action a'_1 such that player 2 is not best-replying at (a'_1, a_2). And player 1 is not best replying at a (since a is not a PNE), so w.p.p. $h^U(a) = (a'_1, a_2)$. Thus some state in which player 2 is not best-replying is reachable from a.

Let $x = (\alpha, \beta)$ be such a state and consider the number of distinct states reachable from x. Player 2 might play any of its actions, so there are at least the l possibilities $(\alpha, 1), ..., (\alpha, l)$ for $h^U(x)$. Since player 1 has no dominant action, there is some $\gamma \in \{1, ..., l\}$ such that, letting $y = (\alpha, \gamma)$, $\alpha \notin BR_1(y)$, so player 1 is not best-replying at y. By the same logic, $(1, \gamma), ..., (k, \gamma)$ are possibilities for $h^U(y)$, and there is a $z = (\delta, \gamma)$ such that player 2 is not best-replying at z and $(\delta, 1), ..., (\delta, l)$ are possibilities for $h^U(y)$.

We've shown that $(\alpha, 1), ..., (\alpha, l), (1, \gamma), ..., (k, \gamma), (\delta, 1), ..., (\delta, l)$ are all reachable from x. Suppose that $\alpha = \delta$. Then $y = z$ and neither player is best replying at y, so all of A is reachable from y. Otherwise,

$$\left|\{(\alpha,1),...,(\alpha,l),(1,\gamma),...,(k,\gamma),(\delta,1),...,(\delta,l)\}\right| \geq k+2l-2$$
$$> k+l.$$

Since x and y are both reachable from a, this completes the proof. □

Proof of Lemma 14. Let $A = \{1,2\}\times\{1,2\}\times\{1,...,k_3\}\times\{1,...,k_4\}$, with $k_3, k_4 \geq 2$. By Observation 5, it suffices to show that h does not succeed on $\mathcal{G}(A)$. Let $U = (u_1,u_2,u_3,u_4) \in \mathcal{G}(A)$ be defined as follows. For every $a = (a_1,a_2,a_3,a_4) \in A$,

$$u_1(a) = \begin{cases} 1 & \text{if } a_1 = a_2 \\ 0 & \text{otherwise,} \end{cases}$$
$$u_2(a) = \begin{cases} 1 & \text{if } a_1 = a_2 \text{ XOR } a_3 = a_4 = 1 \\ 0 & \text{otherwise,} \end{cases}$$
$$u_3(a) = u_4(a) = \begin{cases} 1 & \text{if } a_3 = a_4 \\ 0 & \text{otherwise.} \end{cases}$$

Informally, player 1 always wants to match player 2's action, players 3 and 4 always want to match each other's actions, and player 2 wants to match player 1's action except when players 3 and 4 are both playing 1, in which case player 2 wants to *mismatch* player 1's action.

U has the unique PNE $(2,2,2,2)$. Let $a \in A$ such that $a_3 = a_4$. Then players 3 and 4 are both best-replying, so $h_3^U(a) = h_4^U(a) = 1$. It follows that $(2,2,2,2)$ is not h^U-reachable from $(1,1,1,1)$, so h does not succeed on $\mathcal{G}(A)$. □

Proof of Lemma 15. Let $A = \{1,...,k_1\}\times\{1,...,k_2\}\times\{1,...,k_3\}$, with $k_1,k_2,k_3 \geq 3$. By Observation 5, it suffices to show that h does not succeed on $\mathcal{G}(A)$. Hart and Mas-Colell [7] give an example of a 3-by-3-by-3 generic game on which no historyless uncoupled strategy mapping can succeed. The game is $U = (u_1,u_2,u_3) \in \mathcal{U}(\{1,2,3\}^3)$ where $u_i((x,y,z))$ is the ith coordinate of $M_x[y,z]$, for

$$M_1 = \begin{bmatrix} 0,0,0 & 0,4,4 & 2,1,2 \\ 4,4,0 & 4,0,4 & 3,1,3 \\ 1,2,3 & 1,3,3 & 0,0,0 \end{bmatrix}$$
$$M_2 = \begin{bmatrix} 4,0,4 & 4,4,0 & 3,1,3 \\ 0,4,4 & 0,0,0 & 2,1,2 \\ 1,3,3 & 1,2,2 & 0,0,0 \end{bmatrix}$$
$$M_3 = \begin{bmatrix} 2,2,1 & 3,3,1 & 0,0,0 \\ 3,3,1 & 2,2,1 & 0,0,0 \\ 0,0,0 & 0,0,0 & 6,6,6 \end{bmatrix}.$$

They observe that U has the unique PNE $(3,3,3)$, and prove that if $a \in A$ contains both a 1 and a 2, then for any uncoupled historyless strategy mapping

f, $f^U(a)$ also contains both a 1 and a 2. To prove the lemma, we pad the game with extra actions and show that the expanded game retains this property.

We define the expanded game $U' = (u'_1, u'_2, u'_3) \in \mathcal{G}(A)$ by, for each player i and profile $a = (a_1, a_2, a_3) \in A$,

$$u'_i((a_1, a_2, a_3)) = \begin{cases} 0 & \text{if } a_i > 3 \\ u_i((\min\{a_1, 3\}, \min\{a_2, 3\}, \min\{a_3, 3\})) & \text{otherwise.} \end{cases}$$

So for each new action $a_i > 3$ for player i, a_i is weakly dominated and both other players are indifferent to whether i plays a_i or 3.

Suppose that at $a \in A$ at least one player is playing 1 and at least one player is playing 2. If $a \in \{1, 2, 3\}^3$, then since all the new actions are weakly dominated, Hart and Mas-Colell's analysis applies directly: a player playing 1 and a player playing 2 are best-replying, so $h^U(a)$ contains both a 1 and a 2. Otherwise, one player i is playing $a_i > 3$. In this case the other two players are best-replying, and they played 1 and 2, so $h^U(a)$ again contains both a 1 and a 2. It follows that the players will never reach the PNE $(3, 3, 3)$ starting from, for example, $(1, 2, 1)$, when following h^U. □

Proof of Theorem 18. Except when $A = A_1 \times A_2$ and either $|A_1|$ or $|A_2|$ is 2, this follows directly from Theorem 7. So let $k \geq 2$ and $A = \{1, 2\} \times \{1, ..., k\}$, and assume that some deterministic historyless uncoupled strategy mapping f succeeds on $\mathcal{U}(A)$.

Consider the game $U = (u_1, u_2) \in \mathcal{U}(A)$ defined by

$$u_1(a) = \begin{cases} 1 & \text{if } a_1 = 1 \\ 0 & \text{if } a_1 = 2 \end{cases}$$

$$u_2(a) = \begin{cases} 1 & \text{if } a_2 = a_1 = 1 \text{ or } a_2 \geq a_1 = 2 \\ 0 & \text{otherwise,} \end{cases}$$

for every $a = (a_1, a_2) \in A$. The unique PNE of this game is $p = (1, 1)$, so since we assumed that f succeeds on $\mathcal{U}(A)$, p is f^U-reachable from every $a \in A$.

Define a new game $U' = (u'_1, u'_2) \in \mathcal{U}(A)$ by

$$u'_1(b) = \begin{cases} 2 & \text{if } b_2 \geq x_1 = 2 \text{ and } f_2^U(1, b_2) = 1 \\ u_1(b) & \text{otherwise} \end{cases}$$

$$u'_2(b) = u_2(b),$$

for every $b = (b_1, b_2) \in A$. Informally, each player's preferences are exactly the same as in U, except that player 1 now prefers to play 2 whenever f^U would instruct player 2 to play 1. Notice that U' also has $p = (1, 1)$ as its unique PNE, and that by uncoupledness,$f_2^{U'}(b) = f_2^U(b)$ for every $b \in A$.

Let $a = (1, \alpha) \in A$, for some $\alpha \neq 1$. Notice that $u'_1(a) = u_1(a) = 1$, and consider two cases.

1. $f_2^U(a) = 1$. Then $u'_1((2, \alpha)) = 2$, so player 1 is not U'-best-replying at a. Thus by Observation 4 $f_1^{U'}(a) \neq 1$. Since $f_2^{U'}(a) = f_2^U(a) = 1$, we have $f^{U'}(a) = (2, 1)$.

2. $f_2^U(a) \neq 1$. Then $u_1'((2,\alpha)) = u_1((2,\alpha)) = 0$, so player 1 is U'-best-replying at a, so by Observation 4, $f_1^{U'}(a) = 1$. Since $f_2^{U'}(a) = f_2^U(a) \neq 1$, we have $f^{U'}(a) = (1,\beta)$ for some $\beta \neq 1$.

Now let $b = (2,1)$. Then $u_1'(b) = u_1(b) = 0$, and $u_2'(b) = u_2(b) = 0$, so neither player is best-replying. So by Observation 4 $f_1^{U'}(b) \neq 2$ and $f_2^{U'}(b) \neq 1$, i.e., $f^{U'}(a) = (1,\beta)$ for some $\beta \neq 1$. It follows that $p = (1,1)$ is not $f^{U'}$-reachable from $(2,1)$, so f does not guarantee convergence to a PNE in U', hence f does not succeed on $\mathcal{U}(A)$. $\square$