

On the Structure of Weakly Acyclic Games^{*}

Alex Fabrikant^{1**}, Aaron D. Jagard^{2***}, and Michael Schapira^{3†}

¹ Princeton University, Princeton, NJ, USA
alex@cal.berkeley.edu

² Colgate University, Hamilton, NY, USA & Rutgers University, Piscataway, NJ, USA
adj@dimacs.rutgers.edu

³ Yale University, New Haven, CT, USA & UC Berkeley, Berkeley, CA, USA
michael.schapira@yale.edu

Abstract. The class of *weakly acyclic games*, which includes potential games and dominance-solvable games, captures many practical application domains. Informally, a weakly acyclic game is one where natural distributed dynamics, such as better-response dynamics, cannot enter *in-escapable oscillations*. We establish a novel link between such games and the existence of pure Nash equilibria in subgames. Specifically, we show that the existence of a *unique* pure Nash equilibrium in every *subgame* implies the weak acyclicity of a game. In contrast, the possible existence of *multiple* pure Nash equilibria in every subgame is *insufficient* for weak acyclicity.

1 Introduction

In many domains, convergence to a pure Nash equilibrium is a fundamental problem. In many engineered agent-driven systems that fare best when steady at a pure Nash equilibrium, convergence to equilibrium is expected [7, 9] to happen via *better-response (best-response) dynamics*: Start at some strategy profile. Players take turns, in some arbitrary order, with each player making a better response (best response) to the strategies of the other players, i.e., choosing a strategy that increases (maximizes) their utility, given the current strategies of the other players. Repeat this process until no player wants to switch to a different strategy, at which point we reach a pure Nash equilibrium.

For better-response dynamics to converge to a pure Nash equilibrium regardless of the initial strategy profile, a *necessary* condition is that, from every strategy profile, there exist *some* better-response improvement path (that is, a sequence of players' better responses) leading from that strategy profile to a pure Nash equilibrium. Games for which this property holds are called “weakly

^{*} Partially supported by the DIMACS Special Focus on Communication Security and Information Privacy.

^{**} Supported by a Cisco URP grant and a Princeton University postdoctoral fellowship.

^{***} Partially supported by NSF grants 0751674 and 0753492.

[†] Supported by NSF grant 0331548.

acyclic games” [10, 17]⁴. Both potential games [12, 15] and dominance-solvable games [13] are special cases of weakly acyclic games.

In games that are not weakly acyclic, under better-/best-response dynamics, there are starting states from where the game is guaranteed to oscillate indefinitely. Moreover, the weak acyclicity of a game implies that natural decentralized dynamics (e.g., randomized better-/best-response, or no-regret dynamics) are stochastically guaranteed to reach a pure Nash equilibrium [8, 17]. Thus, weakly acyclic games capture the possibility of reaching pure Nash equilibria via simple, local, globally-asynchronous interactions between strategic agents, independently of the starting state. We assert this is *the* realistic notion of “convergence” in most distributed systems.

1.1 A Motivating Example

We now look at an example inspired by interdomain routing that has this natural form of convergence despite it being, formally, possible that the network will never converge. In keeping with results that we study here, we consider best-response dynamics of a routing model in which each node can see each other node’s current strategy, i.e., its “next hop” (the node to which it forwards its data en route to the destination), as contrasted with models where nodes depending on path announcements to learn this information. (Levin et al. [7] formalized routing dynamics in which nodes learn about forwarding through path announcements.)

Consider the network on four nodes shown in Fig. 1. Each of the nodes 1, 2, and 3 is trying to get a path for network traffic to the destination node d . A strategy of a node i is a choice of a neighbor to whom i will forward traffic; the strategy space of node i , S_i , is its neighborhood in the graph. The utility of d is independent of the outcome, and the utility u_i of node $i \neq d$ depends only on the path that i ’s traffic takes to the destination (and is $-\infty$ if there is no path). We only need to consider the relationships between the values of u_i on all possible paths; the actual values of the utilities do not make a difference. Using $132d$ to denote the path from 1 to 2 to 3 to d , and similarly for other paths, here we assume the following: $u_1(132d) > u_1(1d) > u_1(13d) > -\infty$; $u_2(213d) > u_2(2d) > u_2(21d) > -\infty$; $u_3(321d) > u_3(3d) > u_3(32d) > -\infty$; and $u_i(P) = -\infty$ for all other paths P , e.g., $u_1(12d) = -\infty$. These preferences are indicated by the lists of paths in order of decreasing preference next to the nodes in Fig. 1.




Fig. 1. Instance of the interdomain routing game that is weakly acyclic and has a best-response cycle.

⁴ In some of the economics literature, the terms “weak finite-improvement path property” (weak FIP) and “weak finite best-response path property” (weak FBRP) are also used, for weak acyclicity under better- and best-response dynamics, respectively.

(d, d, d) is a the unique pure Nash equilibrium in this game, and, ideally, the dynamics would always converge to it. However, there exists a best-response cycle:

$$\begin{array}{ccccccc}
 1d & & 132d & & 13d & & 13d & & 1d & & 1d & & 1d \\
 2d & \xrightarrow{1} & 2d & \xrightarrow{3} & 2d & \xrightarrow{2} & 213d & \xrightarrow{1} & 21d & \xrightarrow{3} & 21d & \xrightarrow{2} & 2d & \xrightarrow{1} & \dots \\
 32d & & 32d & & 3d & & 3d & & 3d & & 321d & & 32d
 \end{array}$$

Here, each triple lists the paths that nodes 1, 2, and 3 get; the nodes’ strategies correspond to the second node in their respective paths. The node above the arrow between two triples is the one that makes a best response to get from one triple to the next.

Once the network is in one of these states,⁵ there is a fair activation sequence (i.e., in which every node is activated infinitely often) such that each activated node best responds to the then-current choices of the other nodes and such that the network never converges to a stable routing tree (a pure Nash equilibrium).

Although this cycle seems to suggest that the network in Fig. 1 would be operationally troublesome, it is not as problematic as we might fear. From every point in the state space, there is a sequence of best-response moves that leads to the unique pure Nash equilibrium. We may see this by inspection in this case, but this example also satisfies the hypotheses of our main theorem below. So long as each node has some positive probability of being the next activated node, then, with probability 1, the network will *eventually* converge to the unique stable routing tree, regardless of the initial configuration of the network.

1.2 Our Results

Weak acyclicity is connected to the study of the computational properties of *sink equilibria* [2, 4], minimal collections of states from which best-response dynamics cannot escape: a game is weakly acyclic if and only if all sinks are “singletons”, that is, pure Nash equilibria. Unfortunately, Mirrokni and Skopalik [11] found that reliably checking weak acyclicity is extremely computationally intractable in the worst case (PSPACE-complete) even in succinctly-described games. This means, inter alia, that not only can we not hope to consistently check games in these categories for weak acyclicity, but we cannot even hope to have general short “proofs” of weak acyclicity, which, once somehow found, could be tractably checked.

With little hope of finding robust, effective ways to consistently check weak acyclicity, we instead set out to find *sufficient* conditions for weak acyclicity:

⁵ For example, this might happen if the link between 2 and d temporarily fails. 2 would always choose to send traffic to 1 (if anywhere); 1 would eventually converge to sending traffic directly to d (with 2 sending its traffic to 1), and 3 would then be able to send its traffic along $321d$. Once the failed link between 2 and d is restored, 2’s best response to the choices of the other nodes is to send its traffic directly to d , resulting in the first configuration of the cycle above.

finding usable properties that imply weak acyclicity may yield better insights into at least *some* cases where we need weak acyclicity for the application.

In this work, we focus on general normal-form games. Potential games, the much better understood subcategory of weakly acyclic games, are known to have the following property, which we will refer to as *subgame stability*, abbreviated **SS**: not only does a pure Nash equilibrium exist in the game, but a pure Nash equilibrium exists in each of its *subgames*, i.e., in each game obtained from the original game by the removal of players' strategies. Subgame stability is a useful property in many contexts. For example, in network routing games, subgame stability corresponds to the important requirement that there be a stable routing state even in the presence of arbitrary network malfunctions [5]. We ask the following natural question: When is the strong property of subgame stability *sufficient* for weak acyclicity?

Yamamori and Takahashi [16] prove the following two results⁶:

Theorem: [16] *In 2-player games, subgame stability implies weak acyclicity, even under best response.*

Theorem: [16] *There exist $3 \times 3 \times 3$ games for which subgame stability holds that are not weakly acyclic under best response.*

Thus, subgame stability is sufficient for weak acyclicity in 2-player games, yet is not always sufficient for weak acyclicity in games with $n > 2$ players. Our goal in this work is to (1) identify sufficient conditions for weak acyclicity in the general n -player case; and (2) pursue a detailed characterization of the boundary between games for which subgame stability does imply weak acyclicity and games for which it does not.

Our main result for n -player games shows that a constraint stronger than **SS**, that we term “*unique subgame stability*” (**USS**), is sufficient for weak acyclicity:

Theorem: *If every subgame of a game Γ has a unique pure Nash equilibrium then Γ is weakly acyclic, even under best response.*

This result casts an interesting contrast against the negative result in [16]: *unique* equilibria in subgames guarantee weak acyclicity, but the existence of *more* pure Nash equilibria in subgames can lead to violations of weak acyclicity. Hence, perhaps counter-intuitively, too many stable states can potentially result in persistent instability of local dynamics.

We consider **SS** games, **USS** games, and also the class of *strict and subgame stable* games **SSS**, i.e., subgame stable games which have no ties in the utility functions. We observe that these three classes of games form the hierarchy $\text{USS} \subset \text{SSS} \subset \text{SS}$. We examine the number of players, number of strategies, and the *strictness* of the game (the constraint that there are no ties in the utility function), and give a complete characterization of the weak acyclicity implications of each of these. Our contributions are summarized in Table 1.

⁶ Yamamori and Takahashi use the terms *quasi-acyclicity* for weak acyclicity under best response, and *Pure Nash Equilibrium Property (PNEP)* for subgame stability.

	2 players		3 players				4+ players
	$2 \times M$	$3 \times M$	$2 \times 2 \times 2$	$\frac{2 \times 2 \times M}{2 \times 3 \times M}$	$2 \times 4 \times 4$	$3 \times 3 \times 3$	$2 \times 2 \times 2 \times 2$
\exists pNE	\checkmark (Lma 4)	\nexists (easy)	\checkmark^* (Lma 5)	\nexists (easy)			
SS	\checkmark [16]		\nexists (Thm 2)			\nexists (Thm 2 & [16])	\nexists (Thm 2)
SSS	\checkmark [16]		\checkmark (Lma 5)	\checkmark (Thm 3)	\nexists (Thm 4)	\nexists (Thm 4 & [16])	\nexists (Thm 5)
USS	\checkmark [16]		\checkmark (Thm 1)				

Table 1. Results summary: The impact of USS/SSS/SS on weak acyclicity: \checkmark marks classes with guaranteed weak acyclicity, even under best response; \nexists marks classes which admit counter-examples which are not weakly acyclic even under better response. *: only for strict games

1.3 Other Related Work

Weak acyclicity has been specifically addressed in a handful of specially-structured games: in an applied setting, BGP with backup routing [1], in a game-theoretical setting, games with “strategic complementarities” [3, 6] (a supermodularity condition on lattice-structured strategy sets), and in an algorithmic setting, in several kinds of succinct games [11]. Milchtaich [10] studied Rosenthal’s congestion games [15] and proved that, in interesting cases, such games are weakly acyclic even if the payoff functions (utilities) are not universal but player-specific. Marden et al. [9] formulated the cooperative-control-theoretic consensus problem as a potential game (implying that it is weakly acyclic); they also defined and investigated a time-varying version of weak acyclicity.

1.4 Outline of Paper

In the following, we recall the relevant concepts and definitions in Section 2, present our sufficient condition for weak acyclicity in Section 3, and our characterization of weak acyclicity implications in Section 4.

2 Weakly acyclic games and subgame stability

We use standard game-theoretic notation. Let Γ be a *normal-form game* with n players $1, \dots, n$. We denote by S_i be the *strategy space* of the i^{th} player. Let $S = S_1 \times \dots \times S_n$, and let $S_{-i} = S_1 \times \dots \times S_{i-1} \times S_{i+1} \times \dots \times S_n$ be the cartesian product of all strategy spaces but S_i . Each player i has a *utility function* u_i that specifies i ’s *payoff* in every strategy-profile of the players. For each strategy $s_i \in S_i$, and every $(n-1)$ -tuple of strategies $s_{-i} \in S_{-i}$, we denote by $u_i(s_i, s_{-i})$ the utility of the strategy profile in which player i plays s_i and all other players play their strategies in s_{-i} . We will make use of the following definitions.

Definition 1 (better-response strategies). A strategy $s'_i \in S_i$ is a better response of player i to a strategy profile (s_i, s_{-i}) if $u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i})$.

Definition 2 (best-response strategies). A strategy $s_i \in S_i$ is a best response of player i to a strategy profile $s_{-i} \in S_{-i}$ of the other players if $s_i \in \operatorname{argmax}_{s'_i \in S_i} u_i(s'_i, s_{-i})$

Definition 3 (pure Nash equilibria). A strategy profile s is a pure Nash equilibrium if, for every player i , s_i is a best response of i to s_{-i}

Definition 4 (better- and best-response improvement paths). A better-response (best-response) improvement path in a game Γ is a sequence of strategy profiles s^1, \dots, s^k such that for every $j \in [k - 1]$ (1) s^j and s^{j+1} only differ in the strategy of a single player i and (2) i 's strategy in s^{j+1} is a better response to s_{-i}^j (best response to s_{-i}^j and $u_i(s_i^{j+1}, s_{-i}^j) > u_i(s_i^j, s_{-i}^j)$). The better-response dynamics (best-response dynamics) graph for Γ is the graph on the strategy profiles in Γ whose edges are the better-response (best-response) improvement paths of length 1.

We will use $\Delta R_\Gamma(s)$ and $BR_\Gamma(s)$ to denote the set of all states reachable by, respectively, better and best responses when starting from s in Γ .

We are now ready to define weakly acyclic games [17]. Informally, a game is weakly acyclic if a pure Nash equilibrium can be reached from any initial strategy profile via a better-response improvement path.

Definition 5 (weakly acyclic games). A game Γ is weakly acyclic if, from every strategy profile s , there is a better-response improvement path $s^1 \dots, s^k$ such that $s^1 = s$, and s^k is a pure Nash equilibrium in Γ . (I.e., for each s , there's a pure Nash equilibrium in $\Delta R_\Gamma(s)$.)

We also coin a parallel definition based on best-response dynamics.

Definition 6 (weak acyclicity under best response). A game Γ is weakly acyclic under best response if, from every strategy profile s , there is a best-response improvement path $s^1 \dots, s^k$ such that $s^1 = s$ and s^k is a pure Nash equilibrium in Γ . (I.e., for each s , there's a pure Nash equilibrium in $BR_\Gamma(s)$.)

Weak acyclicity of either kind is equivalent to requiring that, under the respective dynamics, the game has no “non-trivial” sink equilibria [2, 4], i.e., sink equilibria containing more than one strategy profile. Conventionally, sink equilibria are defined with respect to best-response dynamics, but the original definition by Goemans et al. [4] takes into account better-response dynamics as well.

The following follows easily from definitions:

Claim. If a game is weakly acyclic under best response then it is weakly acyclic.

On the other hand, the game in Figure 2, mentioned, e.g., in [8], is weakly acyclic, but not weakly acyclic under best response.

Curiously, all of our results apply both to weak acyclicity in its conventional better-response sense and to weak acyclicity under best response. Thus, unlike weak acyclicity itself, the conditions presented in this paper are “agnostic” to the better-/best-response distinction (like the notion of pure Nash equilibria itself).

We now present the notion of subgame stability.

Definition 7 (subgames). A subgame of a game Γ is a game Γ' obtained from Γ via the removal of players' strategies.

	H	T	X
H	2,0	0,2	0,0
T	0,2	2,0	0,0
X	0,0	1,0	3,3

	c_0		c_1	
	b_0	b_1	b_0	b_1
a_0	2,2,2	1,2,2	2,1,2	2,2,1
a_1	2,2,1	2,1,2	1,2,2	0,0,0

Fig. 2. Matching pennies with a “better-response” escape route, but a best response persistent cycle. **Fig. 3.** $2 \times 2 \times 2$ subgame-stable game with a non-trivial sink

Definition 8 (subgame stability). Subgame stability is said to hold for a game Γ if every subgame of Γ has a pure Nash equilibrium. We use SS to denote the class of subgame stable games.

Definition 9 (unique subgame stability). Unique subgame stability is said to hold for a game Γ if every subgame of Γ has a unique pure Nash equilibrium. We use USS to denote the class of such games.

We will also consider games in which no player has two or more equally good responses to any fixed set of strategies played by the other players. Following, e.g., [14], we define *strict games* as follows.

Definition 10 (strict game). A game Γ is strict if, for any two distinct strategy profiles $s = (s_1, \dots, s_n)$ and $s' = (s'_1, \dots, s'_n)$ such that there is some $j \in [n]$ for which $s' = (s'_j, s_{-j})$ (i.e., s and s' differ only in j 's strategy), then $u_j(s) \neq u_j(s')$.

Definition 11 (SSS). We use SSS to denote the class of games that are both strict and subgame stable.

It's easy to connect unique subgame stability and strictness. To do so, we use the next definition, which will also play a role in our main proofs.

Definition 12 (subgame spanned by profiles). For game Γ with n players and profiles s^1, \dots, s^k in Γ , the subgame spanned by s^1, \dots, s^k is the subgame Γ' of Γ in which the strategy space for player i is $S'_i = \{s^j_i | 1 \leq j \leq k\}$.

Claim. The categories USS , SSS , and SS form a hierarchy: $USS \subset SSS \subset SS$

Proof. $SSS \subset SS$ by definition. To see that $USS \subset SSS$ observe the following. If a game is not strict, there are $s_j, s'_j \in S_j$ and s_{-j} such that $u_j(s_j, s_{-j}) = u_j(s'_j, s_{-j})$. Both strategy profiles in the subgame spanned by (s_j, s_{-j}) and (s'_j, s_{-j}) are pure Nash equilibria, violating unique subgame stability. \square

3 Sufficient condition for weak acyclicity with n players

When is weak acyclicity guaranteed in n -player games for $n \geq 3$? We prove that the existence of a *unique* pure Nash equilibrium in every subgame implies weak acyclicity. We note that this is not true when subgames can contain multiple pure Nash equilibria [16]. Thus, while at first glance, introducing extra equilibria

might seem like it would make it harder to get “stuck” in a non-trivial component of the state space with no “escape path” to an equilibrium, this intuition is false; allowing extra pure Nash equilibria in subgames actually enables the existence of non-trivial sinks.

Theorem 1. *Every game Γ that has a unique pure Nash equilibrium in every subgame $\Gamma' \subseteq \Gamma$ is weakly acyclic under best-response (as are all of its subgames).*

We shall need the following technical lemma:

Lemma 1. *If s is a strategy profile in Γ , and Γ' is the subgame of Γ spanned by $BR_\Gamma(s)$, then any best-response improvement path s, s^1, \dots, s^k in Γ' that starts at s is also a best-response improvement path in Γ .*

Proof. We proceed by induction on the length of the path. The base case is tautological. Inductively, assume s, \dots, s^i is a best-response improvement path in Γ . The strategy s^{i+1} is a best response to s^i in Γ' by some player j . This guarantees that s^i is not a best response by j to s^i_{-j} in Γ' , let alone in Γ , so $\Gamma' \supseteq BR_\Gamma(s) \supseteq BR_\Gamma(s^i)$ must contain a best-response \hat{s}_j^i to s^i_{-j} in Γ , and since s_j^{i+1} is a best-response in Γ' , we are guaranteed that $u_j(\hat{s}_j^i, s^i_{-j}) = u_j(s^{i+1})$, so s^{i+1} must be a best-response in Γ . \square

We may now prove Theorem 1.

Proof (Proof of Theorem 1). To prove Theorem 1, assume that Γ is a game satisfying the hypotheses of the theorem, and for a subgame $\Delta \subseteq \Gamma$, denote by s_Δ the unique pure Nash equilibrium in Δ . We will proceed by induction up the semilattice of subgames of Γ . The base cases are trivial: any $1 \times \dots \times 1$ subgame is weakly acyclic for lack of any transitions. Suppose that for some subgame Γ' of game Γ we know that every strict subgame $\Gamma'' \subsetneq \Gamma'$ is weakly acyclic.

Suppose that Γ' is not weakly acyclic: it has a state s from which its unique pure Nash equilibrium $s_{\Gamma'}$ cannot be reached by best responses. Let Γ'' be the game spanned by $BR(s)$. Consider the cases of (i) $s_{\Gamma'} \in \Gamma''$ and (ii) $s_{\Gamma'} \notin \Gamma''$:

Case (i): $s_{\Gamma'} \in \Gamma''$. This requires that, for an arbitrary player j with more than 1 strategy in Γ' , there be a best-response improvement path from s to some profile \hat{s} where j plays the same strategy as it does in $s_{\Gamma''}$. Take one such j , and let Γ^j be the subgame of Γ' where j is restricted to playing \hat{s}_j only. Since $s_{\Gamma'}$ is in Γ^j , the inductive hypothesis guarantees a best-response improvement path in Γ^j from \hat{s} to $s_{\Gamma'}$. By construction, that path must only involve best responses by players other than j , who have the same strategy options in Γ^j as they did in Γ' , so that path is also a best-response improvement path in Γ' , assuring a best-response improvement path in Γ' from s to $s_{\Gamma'}$ via \hat{s} .

Case (ii): $s_{\Gamma'} \notin \Gamma''$. Then, Γ'' 's unique pure equilibrium $s_{\Gamma''}$ must be distinct from $s_{\Gamma'}$. Since $s_{\Gamma'}$ is the only pure equilibrium in Γ' , $s_{\Gamma''}$ must have an outgoing best-response edge to some profile \hat{s} in Γ' . But the inductive hypothesis ensures that $s_{\Gamma''} \in BR_{\Gamma''}(s)$; by Lemma 1, $s_{\Gamma''} \in BR_{\Gamma'}(s)$, which then ensures that \hat{s} must also be in $BR_{\Gamma'}(s)$, and hence in Γ'' , so $s_{\Gamma''}$ isn't an equilibrium in Γ'' . \square

4 Characterizing the implications of subgame stability

[16] establishes that in 2-player games, subgame stability implies weak acyclicity, even under best response, yet this is not true in 3x3x3 games. We now present a complete characterization of when subgame stability is sufficient for weak acyclicity, as a function of game size and strictness. Our next result shows that the two-player theorem of [16] is maximal:

Theorem 2. *Subgame stability is not sufficient for weak acyclicity even in non-strict $2 \times 2 \times 2$ games.*

Proof. The game in Fig. 3 can be seen to provide the needed counterexample.

However, if we require the games to be strict, subgame stability turns out to be somewhat useful in 3-player games:

Theorem 3. *In any strict $2 \times 2 \times M$ or $2 \times 3 \times M$ game, subgame stability implies weak acyclicity, even under best response.*

We will first need a couple of technical lemmas:

Lemma 2. *In strict games, neither a pure Nash equilibrium and strategy profiles differing from it in only one player's action can be part of a non-trivial sink of the best-response dynamics.*

Proof. A pure Nash equilibrium always forms a 1-node sink. If the game is strict, profiles differing by one player's action have to give that one player a strictly lower payoff, requiring a best-response transition to the equilibrium's sink. Any node connected to either cannot be in a sink.

Lemma 3. *The profiles of a game that constitute a non-trivial sink of the best-response dynamics cannot be all contained within a subgame which is weakly acyclic under best-response.*

Proof. Consider such a non-trivial sink of game Γ contained in such a subgame Γ' . Take a profile s in the sink, and consider the path $P = \{s = s^0, s^1, \dots, s^k\}$ of Γ' best responses that leads to s^k , an equilibrium of Γ' . This path is guaranteed to exist since Γ' is weakly acyclic under best response. Consider the last profile in P , s^a , such that all profiles on P between s and s^a are in the sink.

If $s^a = s^k$ (i.e. if P is entirely in the sink), there has to be a best response transition in Γ from s^k to some s' , since s^k cannot be an equilibrium of Γ and be in a non-trivial sink. If s' were in Γ' , the transition from s^k to s' would have been a best response in Γ' , too, contradicting s^k being an equilibrium of Γ' — thus, s' is not in Γ' , but is in the sink.

If $s^a \neq s^k$, the transition from s^a to s^{a+1} , by some player i , is a best response in Γ' , but not in Γ . So s_i^a is not i 's best response to s_{-i}^a , and thus there is a best response by i from s^a to some s' in Γ . Since s^a to s^{a+1} is *not* a best response transition by i in Γ , $u_i(s') > u_i(s^{a+1})$, and since s^{a+1} is a best response in Γ' , s' must not be in Γ' — but since s^a is in the sink, so is $s' \notin \Gamma'$.

We now start with the corner cases of 3-player, $2 \times 2 \times 2$ games, and 2-player, $2 \times m$ games, where weak acyclicity requires even less than subgame stability. The former result forms the base case for Theorem 3, and both might also be of independent interest.

Lemma 4. *In any $2 \times m$ game, and if there is a pure Nash equilibrium, the game is weakly acyclic, even under best response.*

Proof. In general $2 \times m$ games with pure Nash equilibrium (s^*, t^*) , a non-trivial best-response sink cannot consist of moves by just one player, so the first player will play both of his strategies somewhere in such a sink, including s^* , so there is some (s^*, t') state in the sink. If t' is not a best response to s^* , there would be a best-response transition to the equilibrium (s^*, t^*) , which couldn't happen in a sink. If t' is a best response to s^* , and s^* is a best response to t' , then (s^*, t') is a Nash equilibrium, which couldn't happen in a sink. Lastly, if t' is a best response to s^* , but s^* is not a best response to t' , there has to be some inbound best-response transition into (s^*, t') from another profile in the sink, and that transition then has to involve a move by player 2, from some other state (s^*, t'') , guaranteeing that t'' is not a best response to s^* . Since t^* has to also be a best response to s^* , there is then a best response transition from (s^*, t'') to the equilibrium (s^*, t^*) , concluding the proof. \square

Lemma 5. *In any strict $2 \times 2 \times 2$ game, if there is a pure Nash equilibrium, the game is weakly acyclic, even under best response.*

Proof. In strict $2 \times 2 \times 2$ games, Lemma 2 leaves 4 other strategy profiles, with the possible best-response transitions forming a star in the underlying undirected graph. Since best-response links are antisymmetric ($s \rightarrow s'$ and $s' \rightarrow s$ cannot both be best-response moves), there can be no cycle among those 4 profiles, and thus no non-trivial sink components. \square

Proof (Theorem 3). We treat the $2 \times 2 \times M$ and $2 \times 3 \times M$ cases separately.

The $2 \times 2 \times M$ case: With Lemma 5 as the base case, assume, inductively, that the $2 \times 2 \times 2$ claim holds for all values of M through some $M' - 1$, and suppose some $2 \times 2 \times M'$ game Γ , with strategy sets $\{a_{0,1}\}, \{b_{0,1}\}$, and $\{c_0, \dots, c_{M'-1}\}$, has a non-trivial best-response sink X . WLOG, let (a_0, b_0, c_0) be an equilibrium of Γ .

Lemma 3 guarantees that X is not contained in the subgame Γ_{-c_0} , where only strategy c_0 is removed, leaving a strict, subgame stable $2 \times 2 \times M' - 1$ game, which is weakly acyclic under best response by the inductive hypothesis. But the only profile using c_0 that is allowed to be in X after applying Lemma 2 is (a_1, b_1, c_0) , from which the same lemma guarantees that only player 3 can make a best-response transition. Thus, it can have no inbound best-response transitions by player 3, leaving no way for it to be reached from the rest of X , which can thus not be a sink.

The $2 \times 3 \times M$ case: The $2 \times 3 \times 2$ case is isomorphic to the above. With that as the base case, assume, inductively, that the $2 \times 3 \times M$ claim holds for all

values of M up to some $M' - 1$, and suppose that a $2 \times 3 \times M'$ game Γ , with strategy sets $\{a_{0,1}\}, \{b_{0,1,2}\}$, and $\{c_{0,\dots,M'-1}\}$ has a non-trivial sink X . Again let (a_0, b_0, c_0) , WLOG, be an equilibrium of Γ . The inductive hypothesis and Lemma 3 guarantee that X spans Γ , and, in particular, has at least one node of form $(*, *, c_0)$, and at least one of form $(*, b_0, *)$.

By Lemma 2, the $(*, *, c_0)$ node has to be one of the two nodes $(a_1, b_{1,2}, c_0)$, and that node cannot have an outbound best response by player 1. To be in a non-trivial sink, it has to have an inbound and an outbound best response, one of which is thus by player 3, and the other by player 2, ensuring that *both* of the two nodes $(a_1, b_{1,2}, c_0)$ are in X . One of those will then be player 2's best response to (a_1, c_0) ; WLOG, let that one be (a_1, b_1, c_0) . Then, the only inbound best response to lead to (a_1, b_2, c_0) is by player 3, and player 3 has to have an outbound best response from (a_1, b_1, c_0) to some (a_1, b_1, c_x) .

From (a_1, b_1, c_x) , if there is an outbound best response by player 2, it cannot be to (a_1, b_2, c_x) : else the subgame with strategies $\{a_1\}$, $\{b_{1,2}\}$, and $\{c_{0,x}\}$ is isomorphic to Matching Pennies. Player 2's best response would thus have to instead be to (a_1, b_0, c_x) . From there, player 1 cannot have an outbound best response by Lemma 2, thus requiring player 3 to have a best response to some (a_1, b_0, c_y) ; from there, too, player 1 cannot have a best response by Lemma 2, requiring a best response by player 2. But then, in the $1 \times 3 \times M'$ subgame formed by removing strategy a_0 , for each of player 2's strategies, player 3's best response is to a profile that has an outbound best response by player 2, which precludes an equilibrium.

Thus, from (a_1, b_1, c_x) , the sole possible outbound best response is by player 1, to (a_0, b_1, c_x) .

Consider now the $2 \times 2 \times M'$ subgame Γ_{-b_0} formed by taking away strategy b_0 , and let s^* be its pure Nash equilibrium. If s^* is of form $(a_0, b_{1,2}, c_0)$, that would require that it be player 1's best response in Γ to $(b_{1,2}, c_0)$, thus putting s^* in the sink, in violation of Lemma 2. s^* also cannot be of form $(a_1, b_{1,2}, c_0)$: otherwise, it is in the sink, and yet the only outbound best responses in Γ must be those not in Γ' , i.e. by player 2 to (a_1, b_0, c_0) , in violation of Lemma 2. Thus, s^* has player 3 playing a strategy other than c_0 , which is its best response to one of $(a_1, b_{1,2})$. Player 3's best response to (a_1, b_2) is c_0 , so that cannot be s^* . Player 3's best response to (a_1, b_1) is c_x , but there is an outbound best response by player 1 to (a_0, b_1, c_x) from there, which is within Γ_{-b_0} . Thus, $s^* = (a_0, b_y, c_z)$, for $y \in \{1, 2\}$.

Then, s^* again cannot be in X , since the sole outbound transition in Γ could only be to $(a_0, b_0, *)$, violating Lemma 2. Neither (a_1, b_y, c_z) nor any $(a_0, b_y, *)$ profile can be in X , either, since their best-response transition to s^* in Γ' would also be a best response in Γ , putting s^* into X . If the one (a_0, b_v, c_z) profile with $0 \neq v \neq y$ were in X , it has a best response to s^* in Γ' , so it has to have a best response by player 2 in Γ — but it can't be to (a_0, b_0, c_z) by Lemma 2, and can't be to s^* since s^* can't be in X . Thus, much like in Lemma 2, no profile differing in at most one player's strategy from s^* can be in X , either.

We can now show that nodes $(a_1, b_{0,v}, c_z)$ are both in X , by an argument symmetrical to that for $(a_1, b_{1,2}, c_0)$. The same argument will yield that either b_v or b_0 is the best response to (a_1, c_z) , and that the other one of the two is a best response by player 3. We finish by analyzing the cartesian product of those two cases, and whether $v \in \{1, 2\}$:

$v = 1, b_v$ is best response: The above argument will require that the best response to (a_1, b_1) by player 3, c_x , be neither c_0 nor c_z . If the outbound best response from (a_1, b_1, c_x) is by player 2, to b_0 or b_2 , then either $\{a_1\} \times \{b_{0,1}\} \times \{c_{z,x}\}$ or $\{a_1\} \times \{b_{1,2}\} \times \{c_{0,x}\}$, respectively, form a subgame isomorphic to matching pennies. On the other hand, suppose the outbound best response from (a_1, b_1, c_x) is by player 1, to (a_0, b_1, c_x) . Since $(a_0, b_0, *)$ nodes and $(a_0, b_2, *)$ nodes may not be in X , the only outbound response from there is by player 3, to some c_w , from which the only outbound best response is by player 1 to (a_1, b_1, c_w) , creating a matching pennies subgame with strategies $\{a_{0,1}\} \times \{b_1\} \times \{c_{w,x}\}$.

$v = 1, b_0$ is best response: Then, c_z has to be the best response to (a_1, b_1) by player 3, requiring that $c_x = c_z$, but it was established above that (a_1, b_1, c_x) cannot have an outbound best response by player 2.

$v = 2, b_v$ is best response: An argument symmetrical to the $v = 1, b = 0$ case will show that c_0 , the requisite best response to (a_1, b_1) cannot have an outbound best response by player 2.

$v = 2, b_0$ is best response: Contradiction, since that would require both c_z and c_0 to be the best response to (a_1, b_2) .

Thus, Γ cannot have a non-trivial best-response sink. □

Theorem 3 is maximal. All bigger sizes of 3-player games admit subgame-stable counter-examples that are not weakly acyclic:

Theorem 4. *In non-degenerate⁷ strict 3-player games, the existence of pure Nash equilibria in every subgame is insufficient to guarantee weak acyclicity, for any game with at least 3 strategies for each player, and any game with at least 4 strategies for 2 of the players.*

Proof. The first half of the theorem follows directly from a specific counterexample game in [16]. There, the strict 3-player, $3 \times 3 \times 3$ game in question is stated to demonstrate that SSS does not imply weak acyclicity under best response. Actually, their very same counter-example is not even weakly acyclic under better response. Here, we examine a $2 \times 4 \times 4$ counter-example to establish the second half of the theorem, and a $3 \times 3 \times 3$ counterexample slightly cleaner than the one in [16], both shown in Figure 4.

In each of these three player games, there is a pure Nash equilibrium in the full game, $s^* = (a_1, b_1, c_1)$ in $\Gamma_{3,3,3}$, and $s^* = (a_3, b_3, c_1)$ in $\Gamma_{4,4,2}$, with utility 5 for each of the players. In both, there is a cycle C , every profile in which differs from s^* in at least 2 players' strategies. Any profile (a_i, b_j, c_k) that's neither s^* nor in C yields utilities (i, j, k) . With utilities in C always in $\{4, 5\}$, there is never

⁷ Each player has 2 or more strategies

	c_0			c_1			c_2		
	b_0	b_1	b_2	b_0	b_1	b_2	b_0	b_1	b_2
a_0	0, 0, 0	5, 5, 4	5, 4, 5	4, 5, 5	0, 1, 1	0, 2, 1	5, 5, 4	5, 4, 5	0, 2, 2
a_1	5, 4, 5	1, 1, 0	4, 5, 5	1, 0, 1	5, 5, 5	1, 2, 1	1, 0, 2	1, 1, 2	1, 2, 2
a_2	4, 5, 5	2, 1, 0	2, 2, 0	5, 5, 4	2, 1, 1	2, 2, 1	2, 0, 2	2, 1, 2	2, 2, 2

	c_0				c_1			
	b_0	b_1	b_2	b_3	b_0	b_1	b_2	b_3
a_0	5, 5, 5	0, 1, 0	0, 2, 0	0, 3, 0	5, 5, 4	0, 1, 1	5, 4, 5	0, 3, 1
a_1	1, 0, 0	1, 1, 0	5, 5, 4	5, 4, 5	1, 0, 1	1, 1, 1	4, 5, 5	1, 3, 1
a_2	2, 0, 0	5, 4, 5	2, 2, 0	4, 5, 5	2, 0, 1	2, 1, 1	2, 2, 1	2, 3, 1
a_3	5, 4, 5	4, 5, 5	3, 2, 0	3, 3, 0	3, 0, 1	3, 1, 1	3, 2, 1	5, 5, 5

Fig. 4. 3-player strict subgame stable games that are not weakly acyclic, even under better-response dynamics

an incentive for anyone to unilaterally leave the cycle C , forming a “sheath” of low-utility states separating C from the rest of the game, particularly s^* . Thus C is a persistent cycle. By construction, the game is strict and at each state in C there is a unique player who has a better response to the current state.

Consider any subgame Γ' of either game. If Γ' contains s^* , s^* is a pure Nash equilibrium of Γ' as well.

Suppose Γ' is not the full game. In the course of cycling through C , each strategy of each player is used at least once. Thus, Γ' cannot contain all of C . If it has at least some states of C , pick one state that is in Γ' , and follow the edges of C until you get to a state whose sole outbound better-response move has been “broken” by the better-response strategy being removed in Γ' . This process will terminate, since C is a simple cycle in Γ that had at least one node missing in Γ' . The sole player that had an incentive to move in that state in Γ now no longer has that option, and if he has any other strategy, the resulting state cannot be in C , since C never uses more than 2 strategies of any player i in combination with any fixed s_{-i} . Thus, any other strategy is not an improvement for that player, either, and this new state is thus a pure Nash equilibrium in Γ' .

Lastly, if Γ' contains neither s^* nor any nodes of C , taking the highest-index strategy for each player yields a profile that has to be a pure Nash equilibrium, since the utilities of non- C , non- s^* profiles are just (i, j, k) .

Thus, every subgame is guaranteed to have a pure Nash equilibrium, and, due to C , both games are not weakly acyclic. The theorem holds for games with more strategies by padding the counter-examples above. \square

With 4 or more players, a more mechanistic approach produces analogous examples even with just 2 strategies per player:

Theorem 5. *In a strict n -player game for an arbitrary $n \geq 4$, the existence of pure Nash equilibria in every subgame is insufficient to guarantee weak acyclicity, even with only 2 strategies per player.*

Proof. For strategy profiles in $\{0, 1\}^n$, using indices mod n , set the utilities to:

$$\mathbf{u}(\mathbf{s}) = \begin{cases} (4, \dots, 4) & \text{at } \mathbf{s} = (1, \dots, 1) \\ (3, \dots, 3, \underset{i^{\text{th}}}{2}, 3, \dots, 3) & \text{when } s_{i-1} = s_i = 1, s_{-(i-1,i)} = 0 \\ (3, \dots, 3, \underset{i+1^{\text{th}}}{2}, 3, \dots, 3) & \text{when } s_i = 1, s_{-i} = 0 \\ \mathbf{s} & \text{else (for the “sheath”).} \end{cases}$$

Similarly to Theorem 4, this plants a global pure Nash equilibrium at $(1, \dots, 1)$, and creates a “fragile” better-response cycle. Here, the cycle alternates between profiles with edit distance $n-1$ and $n-2$ from the global pure Nash equilibrium. At every point of the cycle, the only non-sheath profiles 1 step away are its predecessor and successor on the cycle, so the cycle is persistent. Since each profile with edit distance $n-1$ from the equilibrium is covered, removing any player’s 1 strategy breaks the cycle, thus guaranteeing a pure Nash equilibrium in every subgame by the same reasoning as above. \square

We note that, in the counter-example results — Theorems 2, 4, and 5 — the counter-example games of fixed size easily extend to games with extra strategies for some or all players, or with extra players, by “padding” the added part of the payoff table with negative, unique values that, for the added profiles, make payoffs independent of the other players, such as, e.g., $u_i(\mathbf{s}) = -s_i$. This preserves SS, SSS, and USS properties without changing weak acyclicity. Thus, this completes our classification of weak acyclicity under the three subgame-based properties, as shown in Table 1.

5 Concluding remarks

The connection between weak acyclicity and unique subgame stability that we present is surprising, but not immediately practicable: in most succinct game representations, there is no reason to believe that checking unique subgame stability will be tractable in many general settings. In a complexity-theoretic sense, USS is *closer* to tractability than weak acyclicity: Any reasonable game representation will have some “reasonable” representation of subgames, i.e., one in which *checking* whether a state is a pure Nash equilibrium is tractable, which puts unique subgame stability in a substantially easier complexity class, Π_3P , than the class PSPACE for which weak acyclicity is complete in many games.

We leave open the important question of finding efficient algorithms for checking unique subgame stability, which may well be feasible in particular classes of games. Also open and relevant, of course, is the question of more broadly applicable and tractable conditions for weak acyclicity. In particular, there may well be other levels of the subgame stability hierarchy between SSS and USS that could give us weak acyclicity in broader classes of games.

References

1. R. Engelberg and M. Schapira. Hierarchy of Internet routing games, 2010. 12 pages, in submission.
2. A. Fabrikant and C. H. Papadimitriou. The complexity of game dynamics: BGP oscillations, sink equilibria, and beyond. In *Proc. of SODA*, pages 844–853, 2008.
3. J. W. Friedman and C. Mezzetti. Learning in games by random sampling. *Journal of Economic Theory*, 98:55–84, 2001.
4. M. Goemans, V. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *Proceedings of IEEE FOCS 2005*, pages 142–151.
5. T. G. Griffin, F. B. Shepherd, and G. Wilfong. The stable paths problem and interdomain routing. *IEEE/ACM Trans. Netw.*, 10(2):232–243, 2002.
6. N. S. Kulkushkin, S. Takahashi, and T. Yamamori. Improvement dynamics in games with strategic complementarities. *Intl J Game Theory*, 33(2):229–238, 2005.
7. H. Levin, M. Schapira, and A. Zohar. Interdomain routing and games. In *Proceedings of ACM STOC*, pages 57–66, 2008.
8. J. R. Marden, H. P. Young, G. Arslan, and J. S. Shamma. Payoff-based dynamics in multi-player weakly acyclic games. *SIAM J. on Control and Optimization*, 48:373–396, 2009.
9. J.R. Marden, G. Arslan, and J.S. Shamma. Connections between cooperative control and potential games illustrated on the consensus problem. In *Proceedings of the European Control Conference*, July 2007.
10. I. Milchtaich. Congestion games with player-specific payoff functions. *GEB*, 13:111–124, 1996.
11. V. S. Mirrokni and A. Skopalik. On the complexity of nash dynamics and sink equilibria. In *ACM Conference on Electronic Commerce*, pages 1–10, 2009.
12. D. Monderer and L. S. Shapley. Potential games. *GEB*, 14:124–143, 1996.
13. H. Moulin. Dominance solvable voting schemes. *Econometrica*, 47:1337–1351, 1979.
14. N. Nisan, M. Schapira, and A. Zohar. Asynchronous best-reply dynamics. In *Proceedings of the Workshop on Internet Economics*, pages 531–538, 2008.
15. R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
16. T. Yamamori and S. Takahashi. The pure Nash equilibrium property and the quasi-acyclic condition. *Economics Bulletin*, 3(22):1–6, 2002.
17. H. P. Young. The evolution of conventions. *Econometrica*, 61(1):57–84, 1993.