WORKING PAPERS

N° 1627

March 2025

“Nash equilibria are extremely unstable in most games under the utility-taking gradient dynamics”

Aviad Heifetz and Jorge Peña

Toulouse School of Economics

Nash equilibria are extremely unstable in most games under the utility-taking gradient dynamics

Aviad Heifetz*[1] and Jorge Peña[2, 3]

[1]Department of Management and Economics, Open University of Israel
[2]Department of Social and Behavioral Sciences, Toulouse School of Economics, University Toulouse Capitole
[3]Institute for Advanced Study in Toulouse

March 14, 2025

Abstract

In the standard continuous-time choice-taking gradient dynamics in smooth two-player games, each player implicitly assumes that their opponent momentarily maintains their last choice. Contrastingly, in the *utility-taking* gradient dynamics each player implicitly assumes that their opponent momentarily maintains their utility level, by marginally adjusting their choice to that effect. Somewhat surprisingly, employing a transversality argument we find that, in an open and dense set of smooth games, this dynamics is undefined at Nash equilibria. This occurs because, at a Nash equilibrium, the opponent's indifference curve is not locally a function of one's own strategy, making it impossible to specify an opponent's adjustment that would maintain their utility in response to one's own marginal deviation from Nash behavior. Furthermore, when approaching a Nash equilibrium of such a generic game, the utility-taking gradient dynamics either accelerates without bound towards the equilibrium or diverges away from it with unbounded speed.

Keywords: gradient dynamics

1 Introduction

In interactive systems—whether inert or living—stability holds little value without robustness. This is particularly true in social systems, where Nash equilibrium, the hallmark of stability, is meaningful only if behavior naturally tends back to equilibrium after small perturbations. This motivates the long-standing interest in game theory in understanding adaptive behavior away from equilibrium (Weibull, 1997; Fudenberg and Levine, 1998; Hofbauer and Sigmund, 1998; Sandholm, 2010).

The assumptions underlying adaptive behavior vary along a broad spectrum. At one extreme, sophisticated Bayesian learning (Kalai and Lehrer, 1993) assumes full rationality: each player begins with a belief about the opponent's infinite-horizon, history-dependent strategy,

*Corresponding author: aviadhe@openu.ac.il.

and continuously updates this belief while optimizing their own actions. The computational burden of such reasoning is huge. At the other extreme, simple adaptive heuristics (Young, 2004; Hart and Mas-Colell, 2013) rely on bounded memory and limited knowledge, yet still ensure that Nash equilibria arise as rest points of the underlying dynamical system. The most relaxed assumption in terms of memory is that individuals recall only their most recent interaction. In terms of knowledge, players are often assumed to be aware of the set of choices available to them but, in uncoupled dynamics, they neither know their opponent's utility function nor condition their behavior on it.

A natural heuristic that adheres to these minimal assumptions regarding recall and knowledge is best-response dynamics (see e.g. Fudenberg and Levine, 1998). In the continuous-time version of the best-response dynamics (Elkind et al., 2024), each player implicitly assumes that her opponent will momentarily continue to follow their last observed action, and adjusts her own choice in the direction of the best response to the opponent's action, at a speed that equals the utility gain from switching to that best response. Away from Nash equilibrium, this assumption is repeatedly refuted, as the opponent *does* adjust their choice over time. Similarly, with the gradient-ascent dynamics (Mazumdar et al., 2020) each player marginally adjusts their choice in the direction that maximizes their rate of utility increase, at a speed that equals that marginal increase. Here too, the implicit assumption remains that the opponent keeps their last observed choice fixed. This assumption is, in practice, continuously violated except at the rest points—which include the Nash equilibria.

A key advantage of gradient-ascent is its continuity, while best-response dynamics, by contrast, may be discontinuous at certain choice profiles. Another advantage of gradient-ascent dynamics is its local nature: it requires players to evaluate only the relative merit of nearby choices rather than the full range of available actions, as is the case in the best-response dynamics. This makes gradient ascent particularly relevant for real-time decision-making in multi-agent systems, such as autonomous vehicles. However, both best-response and gradient-ascent remain uncoupled dynamics, making them subject to the impossibility result of Hart and Mas-Colell (2003): namely, that there exist games in which no uncoupled dynamics converges to a Nash equilibrium.

This limitation is one motivation for exploring a coupled alternative to gradient ascent, where instead of assuming that her opponent maintains his last observed choice, a player assumes the opponent maintains his most recent payoff level. Under this "utility-taking" gradient dynamics, each player still locally optimizes their own choice at a speed equal to the rate of their optimal marginal improvement. However, the underlying assumption is different: rather than expecting their opponent's strategy to remain fixed, players expect their opponents to adjust their strategy in a way that preserves their current payoff.

The idea of adhering to current utility aligns with observed human behavior. For example, New York taxi drivers were found to decrease their working hours on rainy days, when demand is higher, because they reach their daily income target more quickly (Camerer et al., 1997; Crawford and Meng, 2011). While a more rational approach would be to seize the opportunity for exceptional profits by working even more hours, the drivers instead use their standard income level as a reference point and stop working once they reach it. Similarly, in games, players might naturally assume that opponents respond to external changes—such as one's own strategic deviation—by instinctively adjusting their choices to maintain their payoff, treating such deviations as exogenous shocks beyond their control.

To formalize this idea, we consider a two-player game with choice sets $x_i \in X_i \subseteq \mathbb{R}$, $i = 1, 2$. The combinations of i's choice $x_i \in X_i$, together with the opponent's choice $x_{-i} \in X_{-i}$, and their corresponding utility level $u_{-i} \in \mathbb{R}$ define the set:

$$\Gamma_i = \{(x_i; x_{-i}, u_{-i}) \in X_i \times (X_{-i} \times \mathbb{R}) : U_{-i}(x_i, x_{-i}) = u_{-i}\}.$$

Under the standard uncoupled continuous-time gradient-ascent (and best-response) dynamics, player i implicitly assumes that a deviation from x_i affects u_{-i} but keeps x_{-i} fixed, resulting in a *vertical move* along Γ_i. In contrast, under the utility-taking dynamics, player i assumes that their deviation leads to a *horizontal move*, where x_{-i} changes but u_{-i} remains constant.

A natural question arises: If the standard assumption of choice hysteresis (where the opponent's choice remains fixed) is naive but remains unrefuted at Nash equilibrium, is the alternative assumption of utility hysteresis (where the opponent maintains their payoff level) similarly valid at equilibrium? Somewhat surprisingly, as we show in Section 2, *this question cannot even be meaningfully posed* in generic smooth games, because the utility-taking gradient dynamics is undefined at Nash equilibria. This occurs because, in such games, the opponent's indifference curve is not a well-defined function of one's own strategy near equilibrium, making it impossible to determine their response. In Section 3, we examine what happens *away* from Nash equilibrium under the utility-taking gradient dynamics. We prove that, in generic smooth games, this dynamics exhibits extreme instability: near Nash equilibria, the speeds of players' strategy adjustments tend to infinity, either accelerating without bound towards equilibrium or diverging away from it with unbounded speed. Section 4 summarizes our findings and briefly discusses broader implications.

2 The utility-taking gradient dynamics is undefined at the Nash equilibria of generic games

We consider a game where each player $i \in \{1, 2\}$ selects a strategy from an open set $X_i \subseteq \mathbb{R}$, and the players' utility functions $U_i : X = X_1 \times X_2 \to \mathbb{R}$ are twice continuously differentiable. Throughout, as it is standard in economics, we denote by $-i$ player i's opponent, i.e., $-i = 2$ if $i = 1$ and $-i = 1$ if $i = 2$. We also write x for $(x_1, x_2) \in X$. With this notation, the standard gradient-ascent dynamics in continuous time (Mazumdar et al., 2020) is defined by the system of differential equations:

$$\dot{x}_i = \frac{dx_i}{dt} = \frac{\partial U_i(x)}{\partial x_i}, \quad i = 1, 2, \tag{1}$$

where each player adjusts their choice in the direction that maximizes their marginal utility increase at a speed equal to the rate of this increase, under the naive "choice-taking" assumption that the opponent's choice remains momentarily fixed. Notably, this assumption is valid at Nash equilibria, which are rest points of the dynamics.

As an alternative, we consider the *utility-taking* gradient dynamics. In this dynamics, each player still updates their choice in the direction that maximizes their marginal utility increase at a speed equal to this rate, but under the assumption that the opponent adjusts their choice to maintain their current utility level. Given this assumption, the marginal change in player i's utility is given by:

$$\frac{dU_i(x)}{dx_i} = \frac{\partial U_i(x)}{\partial x_i} + \frac{\partial U_i(x)}{\partial x_{-i}} \left(-\frac{\frac{\partial U_{-i}(x)}{\partial x_i}}{\frac{\partial U_{-i}(x)}{\partial x_{-i}}} \right), \quad i = 1, 2. \tag{2}$$

The term $-\frac{\partial U_{-i}(x)}{\partial x_i} / \frac{\partial U_{-i}(x)}{\partial x_{-i}}$ is the slope of player $-i$'s indifference curve χ_{-i} with respect to x_i, as attained by taking the derivative with respect to x_i on both sides of $U_{-i}(x_i, \chi_{-i}(x_i)) = c$ (where c is a constant), and solving for $\partial \chi_{-i}(x_i)/\partial x_i$. The resulting system of differential equations describing the utility-taking dynamics is:

$$\mathring{x}_i = \frac{dx_i}{dt} = \frac{dU_i(x)}{dx_i} = \frac{\partial U_i(x)}{\partial x_i} - \frac{\frac{\partial U_i(x)}{\partial x_{-i}} \frac{\partial U_{-i}(x)}{\partial x_i}}{\frac{\partial U_{-i}(x)}{\partial x_{-i}}}, \quad i = 1, 2. \tag{3}$$

A natural question is whether Nash equilibria are rest points of utility-taking gradient dynamics. Surprisingly, this question cannot even be meaningfully posed in generic games, that is, in an open and dense set of games under the Whitney (strong) topology (see e.g., Golubitsky and Guillemin, 1973, p. 42-43), which is generated by neighborhoods of games $U \in \mathcal{G}$ of the form

$$\left\{V \in \mathcal{G} : \left\|j_x^2 V - j_x^2 U\right\| < \delta(x) \quad \forall x \in X\right\}$$

where $\delta : X \to \mathbb{R}$ is continuous, and where the 2-jet extensions $j_x^2 V \in J^2(X, \mathbb{R}^2)$ are defined by

$$j_x^2 V = \left(x, V_i(x), \left(\frac{\partial V_i(x)}{\partial x_i}, \frac{\partial V_i(x)}{\partial x_{-i}}\right), \left(\frac{\partial V_i^2(x)}{\partial x_i^2}, \frac{\partial V_i^2(x)}{\partial x_i \partial x_{-i}}, \frac{\partial V_i^2(x)}{\partial x_{-i}^2}\right)\right)_{i=1,2} \in \mathbb{R}^{14}.$$

Specifically, for Nash equilibria (x^*, u^*) in this open and dense set of smooth games, there do not exist open neighborhoods $N_i(x_i^*) \subseteq X_i$ where the opponent's indifference curves $\chi_{-i}(x_i)$ can be implicitly defined by

$$u_{-i}^* = U_{-i}(x_i, \chi_{-i}(x_i)) \tag{4}$$

such that

$$\chi_{-i}(x_i^*) = x_{-i}^*. \tag{5}$$

Consequently, in such a typical game, one cannot even formulate the question: "At a Nash equilibrium, how would player $-i$ react to maintain their utility in response to a marginal change in player i's choice?" As a result, the utility-taking gradient dynamics (3) is undefined at Nash equilibria. This is formalized in the following theorem.

Theorem 1. (i) In a game with twice continuously differentiable utility functions $U_i : X \to \mathbb{R}$, $i = 1, 2$, if indifference curves are well-defined in a neighborhood around a Nash equilibrium (x^*, u^*), then each U_i must be locally flat at x^*. Consequently, a naive assumption that the other player's utility level remains unchanged in response to marginal changes in one's own choice is valid at that Nash equilibrium.

(ii) However, in the space $\mathcal{G}$ of games characterized by tuples $U = (U_i)_{i=1,2}$ of twice continuously differentiable utility functions endowed with the Whitney topology, for games in an open and dense subset $\mathcal{U} \subseteq \mathcal{G}$, indifference curves *cannot be defined* in any neighborhood around Nash equilibria. As a result, the utility-taking assumption is ill-posed at the Nash equilibria of the generic games $U \in \mathcal{U}$, and the utility-taking gradient dynamics is therefore undefined at their Nash equilibria.

Proof. Let (x^*, u^*) be a Nash equilibrium of a game with twice continuously differentiable utility functions $U_i : X \to \mathbb{R}$, $i = 1, 2$, and assume there exist neighborhoods $N_i(x_i^*) \subseteq X_i$ of x_i^* with an indifference curve

$$\chi_{-i} : N_i(x_i^*) \to X_{-i}$$

satisfying (4) and (5). Fully differentiating (4) at x^* with respect to x_i yields

$$0 = \frac{dU_{-i}\left(x_i^*, \chi_{-i}\left(x_i^*\right)\right)}{dx_i} = \frac{\partial U_{-i}\left(x_i^*, \chi_{-i}\left(x_i^*\right)\right)}{\partial x_i} + \frac{\partial U_{-i}\left(x_i^*, \chi_{-i}\left(x_i^*\right)\right)}{\partial x_{-i}} \frac{\partial \chi_{-i}\left(x_i^*\right)}{\partial x_i} \tag{6}$$

Since by (5)

$$U_{-i}\left(x_i^*, \chi_{-i}\left(x_i^*\right)\right) = U_{-i}\left(x^*\right)$$

and since x^* is a Nash equilibrium, we have

$$\frac{\partial U_{-i}\left(x_i^*, \chi_{-i}\left(x_i^*\right)\right)}{\partial x_{-i}} = \frac{\partial U_{-i}\left(x^*\right)}{\partial x_{-i}} = 0 \tag{7}$$

Substituting (7) into (6) yields

$$\frac{\partial U_{-i}\left(x_i^*, \chi_{-i}\left(x_i^*\right)\right)}{\partial x_i} = 0$$

i.e. U_{-i} is locally flat at the Nash equilibrium (x^*, u^*). Therefore, a naive assumption by any individual i that the other player's utility level is marginally unchanged following marginal changes in i's own choice is confirmed at that Nash equilibrium. This proves (i).

However, the utility functions $U = (U_i)_{i=1,2}$ of a game that are locally flat at a Nash equilibrium satisfy there the four equations in two variables

$$\frac{\partial U_i}{\partial x_k} = 0, \quad i, k = 1, 2. \tag{8}$$

Consider the map F defined on the 1-jet extensions

$$J^1\left(X, \mathbb{R}^2\right) = \left\{ j_x^1 U = \left(x, U_i\left(x\right), \frac{\partial U_i\left(x\right)}{\partial x_k} \right)_{i,k=1,2} : U \in \mathcal{G} \right\}$$

by

$$F\left(j_x^1 U\right) = \left(\frac{\partial U_i\left(x\right)}{\partial x_k} \right)_{i,k=1,2}.$$

The map F is the projection of $J^1\left(X, \mathbb{R}^2\right)$ on its last 4 coordinates, and $F^{-1}(0)$ is the closed submanifold of $J^1\left(X, \mathbb{R}^2\right)$ defined by (8).

Now, for every perturbation direction $\pi = (\pi_{ik})_{i,k=1,2} \in \mathbb{R}^4$ of the right-hand side of (8) there exist perturbed games, namely

$$U_i^{\pi,\varepsilon}\left(x_1, x_2\right) = U_i\left(x_1, x_2\right) + \varepsilon \sum_{k=1}^{2} \pi_{ik} x_k, \quad i = 1, 2$$

satisfying at that Nash equilibrium

$$\frac{\partial}{\partial \varepsilon} \left(\frac{\partial U_i^{\pi,\varepsilon}}{\partial x_k} \right) \Bigg|_{\varepsilon=0, x=x^*} = \pi_{ik}, \quad i, k = 1, 2.$$

This shows that F is transversal to the closed submanifold $F^{-1}(0)$ of $J^1\left(X, \mathbb{R}^2\right)$. Therefore, by the jet transversality theorem (see e.g. Hirsch, 1976, theorem 2.8) there exists an open and dense subset $\mathcal{U} \subseteq \mathcal{G}$ of games whose 1-jet extensions $j^1 U$ are transversal to $F^{-1}(0)$. But since $\dim\left(\mathbb{R}^4\right) > \dim\left(\mathbb{R}^2\right)$, this means that for every such generic game $U \in \mathcal{U}$, the system (8) holds for no $x \in X$. In other words, if x^* is a Nash equilibrium of the game U, and therefore satisfies

$$\frac{\partial U_i(x^*)}{\partial x_i} = 0, \quad i = 1, 2,$$

then for some $i \neq k$ we have $\frac{\partial U_i(x^*)}{\partial x_k} \neq 0$, so U_i is not flat at that Nash equilibrium x^*. By (i) it therefore follows that for games in the open and dense subset $\mathcal{U} \subseteq \mathcal{G}$, indifference curves cannot be defined for all individuals around Nash equilibria. This proves (ii). □

Example 1. Consider the two-player game with strategic complements (which is also a potential game, see Monderer and Shapley, 1996) in which the players' utilities are given by:

$$U_i(x_i, x_{-i}) = x_i \left(1 + \frac{x_{-i}}{4}\right) - \frac{x_i^2}{2}, \quad i, j = 1, 2. \tag{9}$$

The players' reaction curves are the graphs of their best-response functions, which are given by $\rho_i^*(x_{-i}) = 1 + x_{-i}/4$, for $i = 1, 2$. The unique Nash equilibrium is located at the intersection of the reaction curves, i.e., at $(x_1^*, x_2^*) = (4/3, 4/3)$. Fig. 1 illustrates the reaction curves, along with several indifference curves of this game. A key observation is that along player 1's reaction curve (and particularly at the Nash equilibrium) her indifference curves are locally flat. This implies that for no point $(\rho_1^*(x_2), x_2)$ on the reaction curve does there exists a neighborhood $N_2(x_2)$ around x_2 where player 1's indifference curve remains a well-defined function of player 2's choice. Specifically: (i) if player 2 increases his choice above x_2, player 1's indifference curve bifurcates, meaning that, to maintain her utility level, player 1 has multiple responses—she may either increase or decrease her choice from $\rho_1^*(x_2)$; and (ii) if player 2 decreases his choice below x_2, then player 1's indifference curve ceases to exist because player 1's utility necessarily decreases regardless of her action. This demonstrates that the question "How would player 1 adjust at $(\rho_1^*(x_2), x_2)$ to maintain her utility level in response to a marginal change in player 2's choice?" is ill-posed, and the utility-taking gradient dynamics is undefined at the Nash equilibrium.

Example 2. Consider the game with utility functions

$$U_i(x_i, x_{-i}) = -\left(x_i - x_{-i}^3\right)^2, \quad i, j = 1, 2$$

and strategy sets $X_1 = X_2 = (-1, 1)$. The players' reaction curves are given by $\rho_i^*(x_{-i}) = x_{-i}^3$, for $i = 1, 2$, and the unique Nash equilibrium is given by $(x_1^*, x_2^*) = (0, 0)$. Fig. 2 illustrates. Each individual's reaction curve coincides with her highest indifference curve, which is a well-defined function of the other's strategy, in particular at the Nash equilibrium. Both utility functions are flat at the Nash equilibrium, as implied by Theorem 1 (i). The utility-taking gradient dynamics is therefore defined at the Nash equilibrium. However, by part (ii) of Theorem 1, there are arbitrarily small perturbations of the utility functions of this game, such that the utility-taking gradient dynamics is undefined at the Nash equilibrium (or Nash equilibria) of the perturbed game.

3 The utility-taking gradient dynamics around Nash equilibria of generic games

In Theorem 1, we established that at the Nash equilibria of typical games, the utility-taking gradient dynamics is undefined. This arises because players cannot meaningfully assess their

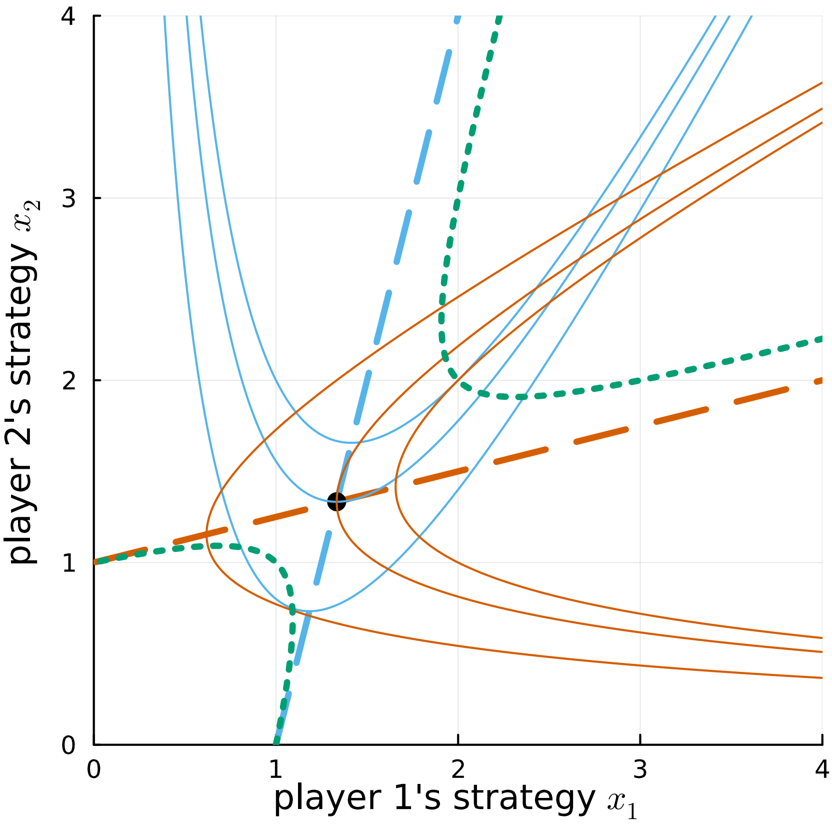

Figure 1: Diagram for the game presented in Example 1. The diagram depicts the reaction curves (*thick, dashed*) together with some indifference curves (*thin, solid*) for player 1 (*blue*) and player 2 (*red*), and tangency points among the player's indifference curves (*dotted, green*). The north-eastern curve of these tangency points consists of the efficient choice profiles, whereas choice profiles on the south-western curve of these tangency points are not efficient. The Nash equilibrium is represented by a *black circle*. The indifference curves are shown for utility values equal to 7/10, 8/9, and 1 for player 1, and 2/3, 8/9, and 1 for player 2.

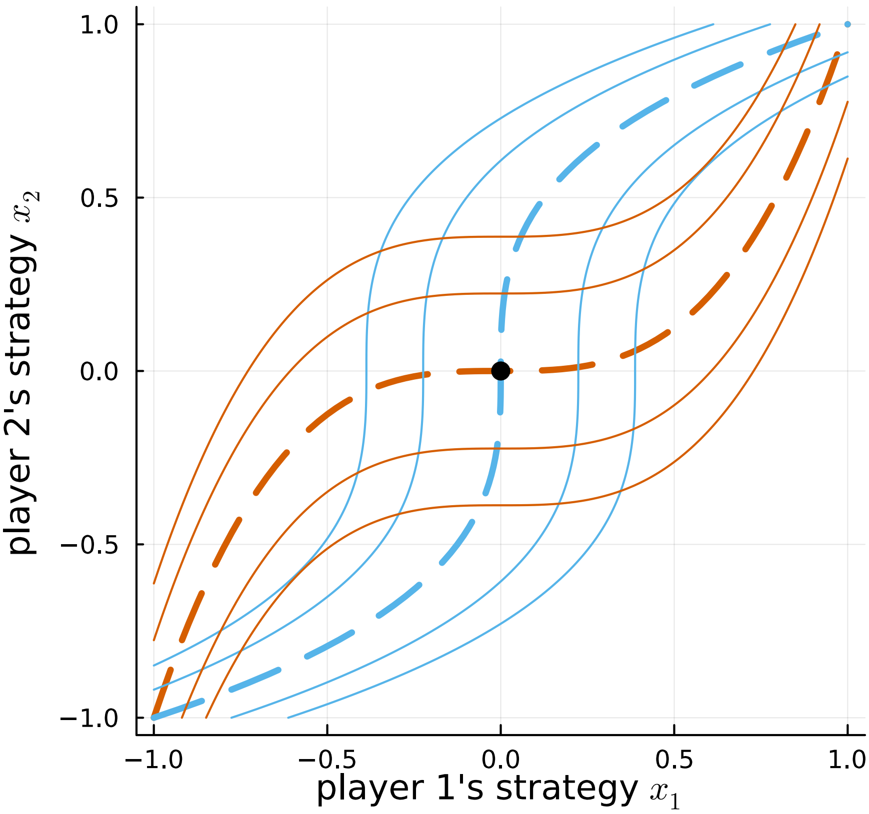

Figure 2: Diagram for the game presented in Example 2. The diagram depicts the reaction curves (*thick, dashed*) together with some indifference curves (*thin, solid*) for player 1 (*blue*) and player 2 (*red*). The Nash equilibrium is represented by a *black circle*. The indifference curves are shown for utility values equal to −0.15, and −0.05 for both players. The reaction curve of each player coincides with their indifference curve for the utility value 0.

opponents' responses to marginal changes in their choices when assuming that their opponents momentarily maintain their utility levels. However, such an assessment *is* possible at choice profiles that lie *away* from the opponent's reaction curve. In these cases, the response is characterized by (2), and the utility-taking gradient dynamics is defined by (3). This leads us to the question: How does this dynamics behave in the vicinity of Nash equilibria? The following theorem addresses this question.

Theorem 2. Consider the space of games where each player's utility function $U_i : X \to \mathbb{R}$, $i = 1, 2$, is twice continuously differentiable. When endowed with the Whitney topology, this space contains an open and dense set of games in which the speed of the utility-taking gradient dynamics (3) tends to infinity as choice profiles approach Nash equilibria.

Proof. Let $U_i : X = X_1 \times X_2 \to \mathbb{R}$, $i = 1, 2$ be twice continuously differentiable utility functions in a game with open strategy sets $X_i \subseteq \mathbb{R}$. As $\partial U_{-i}(x) / \partial x_{-i} = 0$ holds along $-i$'s reaction curve for all $i \in \{1, 2\}$, for a Nash equilibrium x^* we have

$$\lim_{x \to x^*} \frac{\partial U_{-i}(x)}{\partial x_{-i}} = 0.$$

Therefore, if

$$\lim_{x \to x^*} |\mathring{x}_i| = \lim_{x \to x^*} \left| \frac{\partial U_i(x)}{\partial x_i} - \frac{\frac{\partial U_i(x)}{\partial x_{-i}} \frac{\partial U_{-i}(x)}{\partial x_i}}{\frac{\partial U_{-i}(x)}{\partial x_{-i}}} \right| \neq \infty$$

holds, and hence

$$\lim_{x \to x^*} \left| \frac{\frac{\partial U_i(x)}{\partial x_{-i}} \frac{\partial U_{-i}(x)}{\partial x_i}}{\frac{\partial U_{-i}(x)}{\partial x_{-i}}} \right| \neq \infty,$$

then necessarily the numerator of this expression satisfies

$$\lim_{x \to x^*} \frac{\partial U_1(x)}{\partial x_2} \frac{\partial U_2(x)}{\partial x_1} = 0.$$

It then follows that either

$$\frac{\partial U_1(x^*)}{\partial x_2} = 0 \quad \text{or} \quad \frac{\partial U_2(x^*)}{\partial x_1} = 0 \tag{10}$$

must hold. Either of these conditions, together with the first-order conditions satisfied at a Nash equilibrium, constitute, for a given $i \in \{1, 2\}$, a system of equations

$$\begin{pmatrix} \frac{\partial U_i}{\partial x_{-i}} \\ \frac{\partial U_1}{\partial x_1} \\ \frac{\partial U_2}{\partial x_2} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix} \tag{11}$$

with 3 equations in 2 variables. Now, for any perturbation direction

$$\pi = \begin{pmatrix} \pi_{i,-i} \\ \pi_1 \\ \pi_2 \end{pmatrix} \tag{12}$$

of the right hand side of (11), there exist perturbed utility functions, namely

$$\begin{aligned} U_i^{\pi,\varepsilon} &= U_i + \varepsilon \left(\pi_{i,-i} x_{-i} + \pi_i x_i \right) \\ U_{-i}^{\pi,\varepsilon} &= U_{-i} + \varepsilon \pi_{-i} x_k \qquad & i = 1, 2 \end{aligned}$$

of which the ε-derivative of (11) at $\varepsilon = 0$ indeed equals (12). By the jet transversality theorem, and using an argument analogous to that in the proof of Theorem 1(ii), there exists an open and dense set of games $\mathcal{O}_{i,-i}$ in which (11) holds at no $x \in X$, so for the games in the open and dense set $\mathcal{O}_{12} \cup \mathcal{O}_{21}$ the condition (10) holds at no Nash equilibrium. Thus, at every Nash equilibrium x^* of a game in the open and dense set of games $\mathcal{O}_{12} \cup \mathcal{O}_{21}$ it is the case that

$$\lim_{x \to x^*} \left| \frac{\frac{\partial U_i(x)}{\partial x_{-i}} \frac{\partial U_{-i}(x)}{\partial x_i}}{\frac{\partial U_{-i}(x)}{\partial x_{-i}}} \right| = \infty, \quad i = 1, 2,$$

implying that

$$\lim_{x \to x^*} |\mathring{x}_i| = \infty.$$

This proves the result. □

Theorem 2 implies that in sufficiently small neighborhoods of a Nash equilibrium in generic games, the trajectories of the utility-taking gradient dynamics exhibit one of two behaviors. Either they tend towards the Nash equilibrium with unbounded acceleration, effectively 'crashing' onto it, or they diverge away from it, 'exploding' with unbounded speed as the starting choice profile is closer to the equilibrium. Both of these phenomena occur in the game considered by Example 1, which we revisit in the following.

Example 3 (Dynamics for the game of Example 1)**.** The utility-taking gradient dynamics (3) for the game defined by Eq. (9) is given by

$$(\mathring{x}_1, \mathring{x}_2) = \left(\frac{x_2}{4} - x_1 - \frac{1}{16} \frac{x_1 x_2}{\frac{1}{4}x_1 - x_2 + 1} + 1, \frac{x_1}{4} - x_2 - \frac{1}{16} \frac{x_1 x_2}{\frac{1}{4}x_2 - x_1 + 1} + 1 \right). \tag{13}$$

Fig. 3 illustrates the vector field (see `http://bit.ly/3EF4EjY` for an animation of this dynamical system using the vector field explorer *fieldplay*). The Pareto efficient frontier consists of rest points of the dynamical system. When initial conditions are located northeast of the Pareto frontier, the trajectories gradually converge towards it, slowing down as they approach, until ultimately coming to rest on the frontier itself. In contrast, another set of rest points forms a curve of tangency points between the two players' indifference curves, located southwest of the Nash equilibrium. These points are not Pareto efficient. When initial conditions are located southwest of this curve, the trajectories move towards it, decelerating until they settle on the curve. The behavior of the trajectories changes fundamentally for other initial choice profiles. When starting from the northwest or southeast of the Nash equilibrium, trajectories accelerate towards the equilibrium, eventually crashing into one of the best-reply lines at an infinite speed. Conversely, when starting from the northeast or southwest of the Nash equilibrium, trajectories diverge away from the equilibrium while slowing down. Those of these trajectories on or close to the diagonal $x_2 = x_1$ converge towards the nearest rest-point curve, while those further from the diagonal are eventually drawn towards the nearest best-reply line, and accelerate towards it with unbounded speed.

This example highlights a typical case where the Nash equilibrium is highly unstable under the utility-taking gradient dynamics. At the equilibrium itself, the dynamics is undefined. In contrast, the standard (choice-taking) gradient dynamics (1) for this game is given by

$$(\dot{x}_1, \dot{x}_2) = \left(\frac{1}{4}x_2 - x_1 + 1, \frac{1}{4}x_1 - x_2 + 1 \right). \tag{14}$$

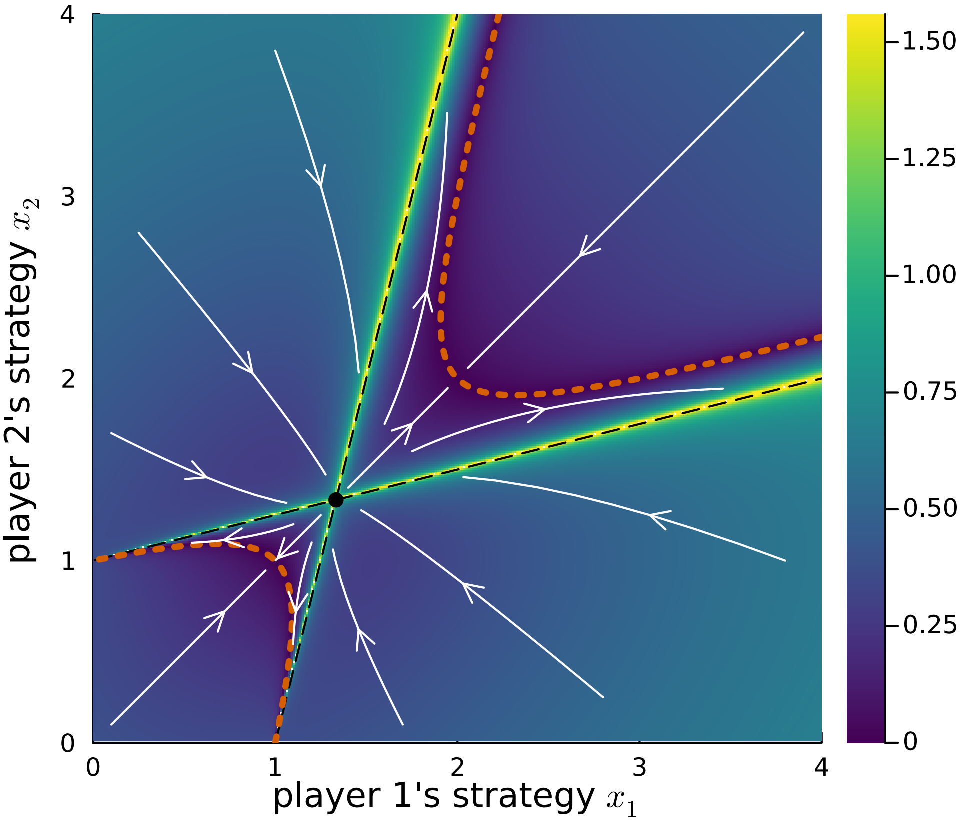

Figure 3: Phase portrait for the utility-taking dynamics (13) of Example 3. Reaction curves (*dashed black lines*), rest points (*dotted orange lines*) together with some trajectories (*solid, white lines*) are shown. The Nash equilibrium is represented by a *black circle*. The heatmap represents the transformed values of the speed of the dynamics, i.e., $z = \sqrt{\mathring{x}_1^2 + \mathring{x}_2^2}$, where non-finite values are ignored. The function output is log-scaled as $\log_{10}(z + 1)$ to enhance visibility. Additionally, to improve contrast, the color range is clipped at the 99.5th percentile of the valid data points.

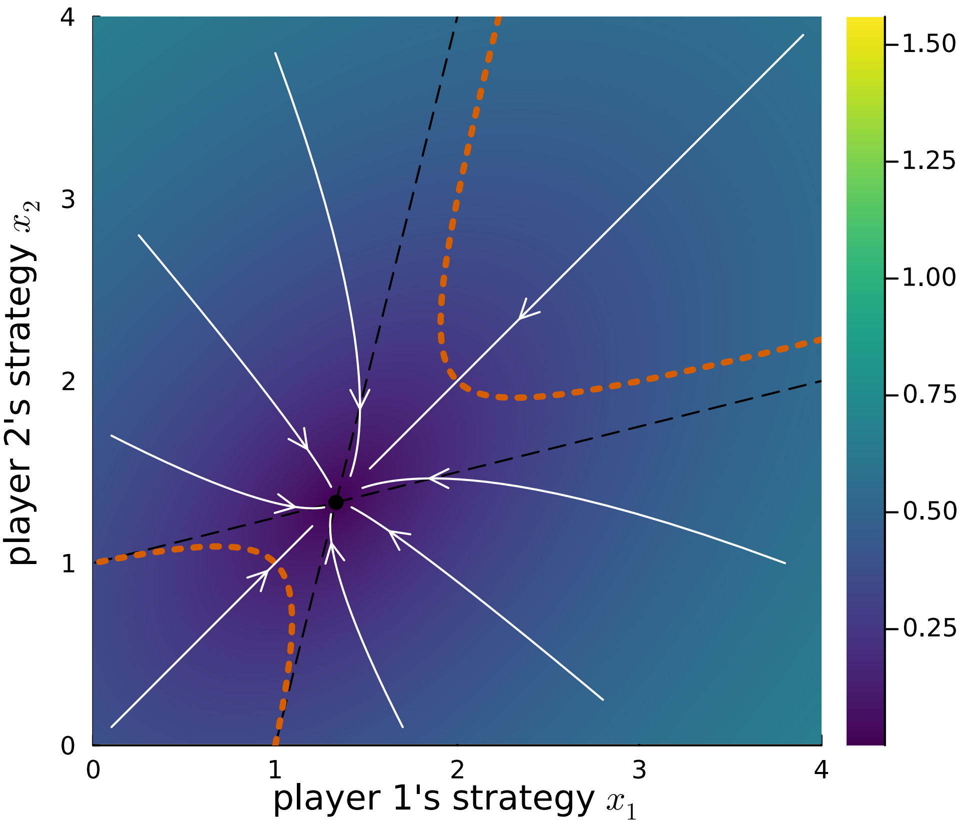

Figure 4: Phase portrait for the choice-taking dynamics (14) of Example 3. Reaction curves (*dashed black lines*), rest points (*dotted orange lines*) together with some trajectories (*solid, white lines*) are shown. The Nash equilibrium is represented by a *black circle*. The heatmap is proportional to the speed of the dynamics using the same transformation as in Fig. 3.

Fig. 4 depicts the vector field of the choice-taking dynamics (see https://bit.ly/4k1PTrB for an animation). It highlights how the two dynamics exhibit significantly different velocities, with trajectories moving in *opposite* directions north-east and south-west of the Nash equilibrium. Unlike the utility-gradient dynamics, the choice-taking gradient dynamics converges globally to the Nash equilibrium, which is a rest point of the dynamical system.

In contrast, we recall that in the atypical game of Example 2, the utility-taking gradient dynamics *was* well-defined at the Nash equilibrium. As it will be seen below, in this atypical game the Nash equilibrium is an asymptotically stable rest-point of the dynamics.

Example 4 (Dynamics for the game of Example 2)**.** The vector field of the utility-taking gradient dynamics (3) is given by

$$(\mathring{x}_1, \mathring{x}_2) = \left(-2\left(1 - 9x_1^2x_2^2\right)\left(x_1 - x_2^3\right), -2\left(1 - 9x_1^2x_2^2\right)\left(x_2 - x_1^3\right)\right). \tag{15}$$

Fig. 5 illustrates the dynamics (see https://bit.ly/3WVwhvh for an animation). In this example, the utility-taking gradient dynamics converges towards the Nash equilibrium $(0, 0)$ from all directions, progressively slowing down until coming to a complete stop at $(0, 0)$. The dynamics remain well-defined at the Nash equilibrium, which serves as a rest point of the dynamics.

For this game, the standard choice-taking gradient dynamics (1) is given by

$$(\dot{x}_1, \dot{x}_2) = \left(-2\left(x_1 - x_2^3\right), -2\left(x_2 - x_1^3\right)\right). \tag{16}$$

Fig. 6 depicts the vector field of the choice-taking dynamics (see https://bit.ly/40V9OzX for an animation). The ratio between the two dynamics

$$\frac{\mathring{x}_1}{\dot{x}_1} = \frac{\mathring{x}_2}{\dot{x}_2} = 1 - 9x_1^2x_2^2$$

is positive in a neighborhood of the Nash equilibrium $(0, 0)$, and tends to 1 as (x_1, x_2) tends to $(0, 0)$. Therefore, in such a neighborhood both dynamics share the same paths, and move in the same direction along each path, with relative speeds tending to 1 (and both tending to 0). However, beyond this neighborhood of the Nash equilibrium $(0, 0)$, i.e., where $1 - 9x_1^2x_2^2 < 0$ holds (in the neighborhoods of X's corners $(1, 1), (1, -1), (-1, 1), (-1, -1)$) the trajectories of the two dynamics proceed in opposite directions, and thus the two dynamics have a very different global behavior.

4 Concluding remarks

Roger Myerson, in his game theory textbook, Myerson, 1991, p. 106 poses the question: *"When asked why players in a game should behave as in some Nash equilibrium, my favorite response is to ask, 'Why not?' and to let the challenger specify what he thinks the players should do."* In this work, we have provided one such specification. We introduced the utility-taking gradient dynamics, for which Myerson's question 'Why not play a Nash equilibrium?' is ill-posed in generic games. The implicit assumption about the opponent in this dynamics is congruent with the natural human instinct to hold on to one's lot by adjusting one's choice to that effect. This human tendency has documented empirical support (Camerer et al., 1997; Crawford and Meng, 2011).

Our challenge to Nash behavior differs from that proposed by Milionis et al. (2023). Their work presents an example of a *non-generic* game in which no continuous-time game

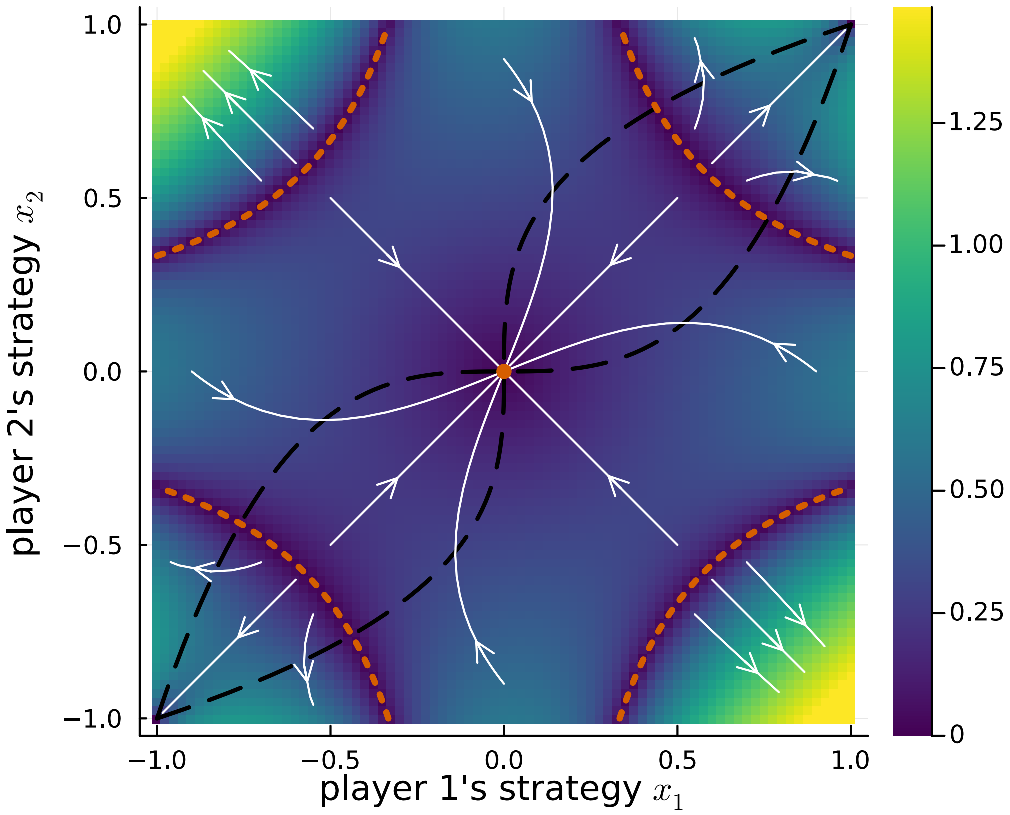

Figure 5: Phase portrait for the utility-taking dynamics (15) of Example 4. Reaction curves (*dashed black lines*), rest points (*dotted orange lines* and *orange circle* together with some trajectories (*solid, white lines*) are shown. The Nash equilibrium is the rest point $(0, 0)$ located at the intersection of the reaction curves. The heatmap represents the transformed values of the speed of the dynamics, i.e., $z = \sqrt{\mathring{x}_1^2 + \mathring{x}_2^2}$, where non-finite values are ignored. The function output is log-scaled as $\log_{10}(z+1)$ to enhance visibility. Additionally, to improve contrast, the color range is clipped at the 99.5th percentile of the valid data points.

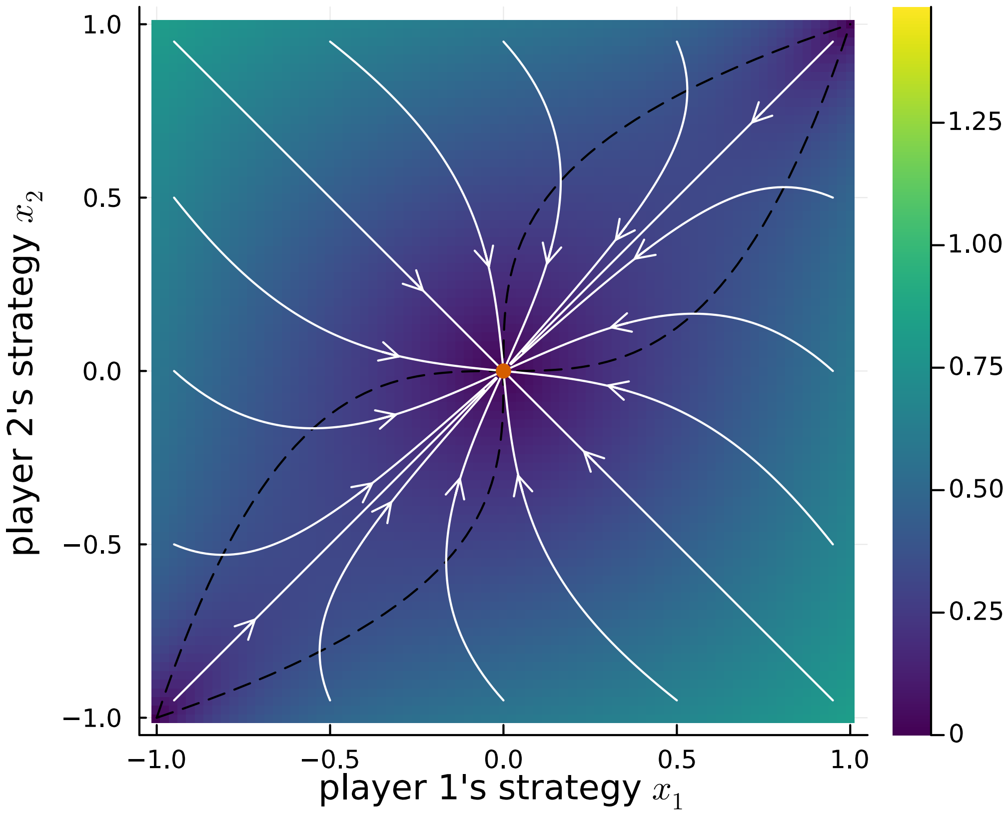

Figure 6: Phase portrait for the choice-taking dynamics (16) of Example 4. Reaction curves (*dashed black lines*), rest point (*orange circle*) together with some trajectories (*solid, white lines*) are shown. The Nash equilibrium is the rest point $(0, 0)$ located at the intersection of the reaction curves. The heatmap is proportional to the speed of the dynamics using the same transformation as in Fig. 5.

dynamics—defined as a dynamic system whose rest points include the set of Nash equilibria—converges to the Nash equilibrium, even when starting from an arbitrarily small neighborhood of it. This critique led Hakim et al. (2024) to favor the limit points of certain well-established game dynamics, such as noisy replicator dynamics, over Nash equilibria.

In contrast, we do not restrict our analysis to dynamics in which Nash equilibria are *a priori* designated as rest points. Instead, we broaden our scope to encompass natural dynamics more generally, where what constitutes 'natural' may depend on the context. The utility-taking assumption explored here may be particularly relevant in settings where payoff levels serve as benchmarks—such as in the anonymous free-market example of taxi drivers' work supply discussed in Camerer et al. (1997). It may also apply in highly personalized, non-anonymous interactions, such as within families, where individuals might primarily respond to one another's well-being rather than to specific strategic choices.

Ratliff et al. (2014) established the converse of Theorem 2 for the standard *choice-taking* gradient dynamics (1), demonstrating the stability of Nash equilibria in an open and dense family of smooth games. From this perspective, Theorem 2 underscores how the stability of equilibria is highly sensitive to the assumptions individuals make about others when employing adaptive heuristics.

We examined two possible naive assumptions a player might make—either assuming that the other player keeps their choice fixed or that they maintain their utility level in response to marginal changes in one's choice. But what if, instead of being naive, individuals could correctly anticipate how others would marginally adjust their choices at any given moment and respond optimally to these correct anticipations at every possible choice profile? This leads to the concept of subgame-perfect feedback equilibria in differential games. However, even in relatively simple differential games with smooth utility functions—such as those with linear-quadratic utility functions—there may exist a continuum of subgame-perfect feedback equilibria (see, e.g., Lambertini, 2018, p. 16 and Proposition 7.4). In other words, assuming full sophistication and perfect foresight for individuals who can continuously adjust their behavior does not necessarily determine how they will or should behave.

Acknowledgements

A. H. is grateful to Herakles Polemarchakis and Enrico Minelli for numerous discussions on related themes, while working on their joint project "Affective interdependence and welfare", and to the Institute for Advanced Study in Toulouse for its hospitality. We acknowledge funding from the French National Research Agency under the Investments for the Future (Investissements d'Avenir) program, grant ANR-17-EURE-0010. We are also grateful to Andrei Kashcha, the developer of *fieldplay* (https://github.com/anvaka/fieldplay), used here under its MIT License to visually explore the vector fields of the dynamical systems we study.

References

Camerer, C., Babcock, L., Loewenstein, G., Thaler, R., 1997. Labor supply of New York City cabdrivers: One day at a time. The Quarterly Journal of Economics 112, 407–441.

Crawford, V.P., Meng, J., 2011. New York City cab drivers' labor supply revisited: Reference-dependent preferences with rational-expectations targets for hours and income. American Economic Review 101, 1912–32.

Elkind, E., Ghosh, A., Goldberg, P.W., 2024. Continuous-time best-response and related dynamics in Tullock contests with convex costs. arXiv preprint arXiv:2402.08541 .

Fudenberg, D., Levine, D.K., 1998. The theory of learning in games. MIT Press.

Golubitsky, M., Guillemin, V., 1973. Stable mappings and their singularities. Springer-Verlag.

Hakim, R., Milionis, J., Papadimitriou, C., Piliouras, G., 2024. Swim till you sink: Computing the limit of a game, in: Schäfer, G., Ventre, C. (Eds.), Algorithmic Game Theory, Springer Nature, Cham. pp. 205–222.

Hart, S., Mas-Colell, A., 2003. Uncoupled dynamics do not lead to Nash equilibrium. American Economic Review 93, 1830–1836.

Hart, S., Mas-Colell, A., 2013. Simple adaptive strategies: from regret-matching to uncoupled dynamics. World Scientific.

Hirsch, M.W., 1976. Differential topology. Springer Science & Business Media.

Hofbauer, J., Sigmund, K., 1998. Evolutionary games and population dynamics. Cambridge University Press.

Kalai, E., Lehrer, E., 1993. Rational learning leads to Nash equilibrium. Econometrica 61, 1019–1045.

Lambertini, L., 2018. Differential games in industrial economics. Cambridge University Press.

Mazumdar, E., Ratliff, L.J., Sastry, S.S., 2020. On gradient-based learning in continuous games. SIAM Journal on Mathematics of Data Science 2, 103–131.

Milionis, J., Papadimitriou, C., Piliouras, G., Spendlove, K., 2023. An impossibility theorem in game dynamics. Proceedings of the National Academy of Sciences 120, e2305349120.

Monderer, D., Shapley, L.S., 1996. Potential games. Games and Economic Behavior 14, 124–143.

Myerson, R.B., 1991. Game theory. Harvard University Press.

Ratliff, L.J., Burden, S.A., Sastry, S.S., 2014. Genericity and structural stability of non-degenerate differential Nash equilibria, in: 2014 American Control Conference, pp. 3990–3995.

Sandholm, W.H., 2010. Population games and evolutionary dynamics. MIT Press.

Weibull, J.W., 1997. Evolutionary game theory. MIT Press.

Young, H.P., 2004. Strategic learning and its limits. Oxford University Press.