Combinator Dynamics: Multi-Dimensional Evolutionary Learning*

Srinivas Arigapudi[†] Omer Edhan[‡] Yuval Heller[§] Ziv Hellman[¶]

June 24, 2026

Abstract

In complex strategic settings with multiple dimensions, players often simplify decisions by breaking them down into separate dimensions. In contrast to standard replicator dynamics, which assume perfect multi-dimensional coordination, our combinator dynamics model captures cases in which players revise actions dimension by dimension, imitate successful traits in each dimension, and then combine the selected traits into a new action. We analyse key properties of these dynamics – monotonicity, stationarity, and stability. We then extend the model to heterogeneous settings, where some players imitate entire multi-dimensional actions while others combine traits

* We thank seminar and conference participants at the LEG conference (Lucca, 2022), Bar-Ilan University, the Hebrew University of Jerusalem, Technion (Israel Institute of Technology), Corvinus University of Budapest, Stony Brook University, UC Santa Barbara, UC San Diego, the University of Manchester, the University of Glasgow, three anonymous referees and the editor, Prof. Juuso Toikka, for various helpful comments. SA gratefully acknowledges financial support from the ANRF ARG MATRICS grant (ANRF/ARGM/2025/000809/QSS) and an Initiation grant at IIT Kanpur. OE gratefully acknowledges research support from the British Academy (R128030). All authors gratefully acknowledge research support from the Israel Science Foundation (#448/22 and #458/25). A previous version of the paper was called "Mentors and Recombinators: Multi-Dimensional Social Learning."

[†]Department of Economic Sciences, IIT Kanpur, India, arigapudi@iitk.ac.in.

[‡]Department of Economics, University of Manchester, UK, omeredhan.idan@manchester.ac.uk

[§]Department of Economics, Bar-Ilan University, Israel, yuval.heller@biu.ac.il.

[¶]Department of Economics, Bar-Ilan University, Israel, ziv.hellman@biu.ac.il

imitated separately in each dimension. We apply the framework to organisational decision-making, where decentralised adoption of practices can generate coordination failures, persistent performance differences, and a trade-off between specialised and versatile organisational traits.

Keywords: multi-dimensional reasoning, non-monotone dynamics, evolutionary stability, imitation. **JEL codes**: C73, D83, L20

1 Introduction

In many strategic settings, the strategy space is large and complex, potentially beyond the capacity of players to formulate strategic choices that take into account the entire strategy space as a single unit. To manage this complexity, players often break down their strategic choices into multiple dimensions, implementing separate decisions for each dimension (Arad and Rubinstein, 2019). This possibility is especially natural when the player is interpreted as an organisation, because complex decisions are often decomposed across subunits, but the same logic can arise in many other settings with multi-dimensional actions. For instance, an early career researcher navigating professional development may consult one mentor for choosing a research area, another for presenting at conferences, and yet another for selecting and collaborating with co-authors.

A widely used model of learning is based on the replicator dynamics (Taylor and Jonker, 1978) from evolutionary game theory. This approach considers a large population of players who are randomly matched to play a symmetric normal-form game. Players periodically revise their strategic choices (henceforth, *actions*). Under replicator dynamics, when an agent updates a multi-dimensional action, he or she aggregates all strategic dimensions perfectly and selects a new action that yields high payoffs—even when that differs across multiple dimensions from the previous action.

We introduce in this paper a novel model for multi-dimensional learning, *combinator dynamics*, in which, in contrast to the replicator model, agents evaluate each dimension separately, selecting from each a *trait* that typically yields high payoffs. These traits, one from each dimension, are then combined to form a new action.

One source of inspiration for our approach is the contrast between biological asexual and sexual reproduction strategies. Asexual reproduction is best modelled by replicator dynamics: under such reproduction, a parent genotype in its entirety is inherited by an offspring. Sexual reproduction in contrast is a combinator form of multi-dimensional learning — genes from more than one parent are combined to form an offspring, with the determination of which gene is inherited from each parent conducted separately at each location in the genotype.

Our main results analyse three key aspects of combinator dynamics: monotonicity, stationarity, and stability. Standard replicator dynamics are payoff-monotone, meaning that an action with a higher payoff grows at a higher rate. We first show that combinator dynamics are payoff-monotone only at the trait level: within each dimension, traits with higher payoffs grow at a higher rate, but this does not necessarily hold for multi-dimensional actions.

A distribution of actions (henceforth, *state*) is stationary if, once reached, the population remains in that state indefinitely. In replicator dynamics, a state is stationary if and only if all incumbent actions yield the same payoff. By contrast, our second result shows that in combinator dynamics, a state is stationary if and only if (1) all traits yield the same payoff, and (2) the frequency of each multi-dimensional action equals the product of the frequencies of its traits.

A stationary state is (asymptotically) stable if any nearby population converges to it over time. In replicator dynamics, this holds if (1) the Jacobian matrix of the payoff function is negative definite over the population's support, and (2) any external action yields a lower payoff than the incumbents. Our third result adapts these conditions to characterise stability under combinator dynamics: (1) the Jacobian matrix of an alternative payoff function—defined as the product of the traits' payoffs and their frequencies—is negative definite, and (2) any external trait invading the population yields a lower payoff than the incumbents, though no such restriction applies to invading actions involving multiple new traits.

The formal results are illustrated by three examples. Example 1 studies complementary improvements: changing either dimension alone is harmful, whereas changing both dimensions together is beneficial. Example 2 studies congestion with cross-dimensional lignment: each trait is valuable when it is relatively scarce, but a full action receives an additional payoff when its traits are aligned. The stable state under combinator dynamics can then contain full actions with different payoffs even though all traits have the same payoff. Example 3 studies specialised and versatile traits: a specialised trait performs well in one configuration and poorly in another, whereas a versatile trait performs moderately well across configurations.

We then extend the model to *recombinator dynamics*. The recombination parameter is a parameter of the revision protocol. With probability $1 - r$, a revising agent copies a whole action, as in the replicator dynamic. With probability r, the revising agent samples traits independently across dimensions and recombines them, as in the combinator dynamic. This extension allows us to study the continuous transition from whole-action imitation to component-wise imitation. The corresponding stability conditions combine restrictions on full actions with restrictions on traits.

Section 6 applies the theory to organisational decision-making. Organisations often delegate the assessment of different dimensions of activity to distinct divisions, each responsible for identifying high–payoff traits within its domain. The three examples then illustrate complementary organisational change, persistent performance differences generated by imperfect fit among components, and a trade-off between specialised and versatile organisational traits.

The paper's main contribution is to introduce a tractable evolutionary dynamic for multi-dimensional environments in which agents revise components of an action separately. Component-wise imitation can make inefficient complementary states locally stable (Example 1), equalise trait payoffs while inducing unequal action payoffs (Example 2), and shape the stability of specialised versus versatile traits (Example 3). The stability analysis refines Arad and Rubinstein's (2019) multi-dimensional equilibrium and yields sharper predictions in games with large action sets. As shown in Section 6, the model can account for several stylised facts in organisational decision-making.

The paper is structured as follows. In Section 2, we discuss the related literature.

Section 3 presents our baseline model of the combinator dynamics. Section 4 presents our results. In Section 5 we extend our model to heterogeneous environments, where some players combine imitated traits, while others imitate whole actions. Section 6 applies the framework to organisational decision-making. We conclude in Section 7. Technical proofs are presented in the Appendix.

2 Related Literature

In this section we review the literature most closely related to the paper: multi-dimensional reasoning in games, evolutionary and imitative dynamics, and multi-dimensional social learning. (We postpone the discussion of the related organisational behaviour literature to Section 6.)

Multi-dimensional reasoning in games Arad and Rubinstein (2012) and Arad and Penczynski (2024) provide experimental evidence that players perceive games with large action sets as multi-dimensional, reasoning about each strategic dimension independently.[1] Motivated by this, Arad and Rubinstein (2019) introduce the (non-evolutionary) solution concept of *multi-dimensional (MD) equilibrium*. An action profile a is an MD equilibrium if any better reply against itself requires changing multiple traits (i.e., if $u(a',a) > u(a,a)$ implies that a' differs from a in at least two dimensions).[2]

For pure strategies, an MD equilibrium is similar to our notion of an asymptotically stable pure state in the combinator dynamics.[3] However, the two concepts diverge significantly in mixed settings, where our solution concept is a refinement of theirs. A mixed MD equilibrium is a distribution of actions where no player can gain by changing a single trait of an incumbent's action. This concept is too weak in social learning

[1] Another interpretation of this behaviour is that players engage in narrow bracketing (Thaler, 1999; Read et al., 2000): they consider each strategic dimension in isolation, adopting a trait for each dimension without assessing the overall consequences of combining traits across dimensions.

[2] When the set of actions is large, Arad and Rubinstein (2019) allow multiple actions to correspond to each vector of traits (in which case, the vector of traits is interpreted as a cell that contains a set of similar strategies). We focus here on the application of their solution concept to "product edited games" (which are the games analysed in the present paper) in which each vector of traits corresponds to a single action.

[3] Specifically, any strict MD equilibrium is an asymptotically stable pure state in the combinator dynamics, and any asymptotically stable pure state in the combinator dynamics is an MD equilibrium.

contexts, as it permits frequencies of actions to deviate significantly from the product of trait frequencies—an impossibility in a stationary state under combinator dynamics, as shown in Proposition 2.

Consider Example 2 in Section 4.2. A state in which half the population chooses one aligned action and the other half chooses the other aligned action constitutes a mixed MD equilibrium. However, it is neither stable nor stationary under the combinator dynamics. When agents revise each dimension separately, some newly formed actions combine a trait from one aligned action with a trait from the other, thereby generating misaligned actions that are absent from the original mixed state. A key contribution of the paper is to show that the evolutionary approach reveals weaknesses of the MD equilibrium concept—at least in environments with social learning—and suggests refinements that yield sharper predictions.

Evolutionary and imitative dynamics The evolutionary game theory literature (pioneered in Maynard Smith and Price, 1973) considers a game that is played over and over again by biologically or socially conditioned players who are randomly drawn from large populations. Occasionally, new agents join the population (or incumbents revise their behaviour), and they learn how to play based on observing the (possibly noisy) behaviour and payoffs of some of the incumbent agents (see, Weibull, 1997 and Sandholm, 2010 for a textbook introduction, and Newton, 2018 for a comprehensive recent survey of the literature).

The commonly-applied *replicator dynamic* (Taylor and Jonker, 1978) models how aggregate behaviour evolves, with the growth rate of agents playing each action proportional to their average payoff. Originally developed for genetic evolution, it has been widely applied to social learning, where new agents tend to imitate successful incumbents (see, e.g., Börgers and Sarin, 1997; Hopkins, 2002; Cressman and Tao, 2014; Sawa and Zusai, 2014; Mertikopoulos and Sandholm, 2018).

The vast majority of the evolutionary game theory literature assumes that the learning process is one-dimensional.[4,5] In particular, most of the literature on imitative processes assumes that a new agent mimics the behaviour of a single incumbent. In what follows, we describe the relatively small literature that deals with multi-dimensional learning, in which new agents may combine the learning of various traits from different incumbents.[6]

Multi-dimensional social learning The multi-dimensional social learning model mirrors sexual inheritance: new agents combine traits from multiple incumbents, akin to how offspring inherit genes from both parents. The stability of phenotypic behaviour that is determined by the combination of genes (at different loci in the DNA) has long been studied in the biological literature (see, e.g., Karlin, 1975; Eshel and Feldman, 1984; Matessi and Di Pasquale, 1996).

Waldman (1994) studies a setting where agents choose two bias levels (e.g., overconfidence and work disutility). He shows that under sexual inheritance, a pair of biases can be evolutionarily stable if each is optimal given the other—termed "second-best adaptations." Frenkel et al. (2018) extend this to a continuum of bias levels, showing that while second-best adaptations disappear, biases that roughly compensate for each other can persist for long periods.[7]

Several studies explore the link between evolutionary dynamics and learning algorithms (Chastain et al., 2014; Barton et al., 2014; Meir and Parkes, 2015). They show that sexual inheritance enables regret minimisation, convergence to Nash equilibria, and can be viewed as a common-interest game among genetic loci. Edhan et al. (2017) highlight

[4] In a recent survey, Bisin and Verdier (2023) point out that most of the literature on the economic theory of cultural transmission is on the evolution of a single cultural trait which excludes cultures as systems of traits whose combinations and interactions are essential to generating new cultural meanings.

[5] Some papers (for example Cressman et al., 2000, Hashimoto, 2006, and Sawa and Zusai, 2019) study models where players engage in multiple games simultaneously, and their overall payoff is determined by the cumulative sum of payoffs obtained in each distinct game. In these studies, each game is conceptualised as a dimension. This contrasts with our approach here, in which a player is associated with multiple traits, each of which may be considered a selection of an action from one set of possible actions out of a profile of action sets, but the payoff for the player may be an arbitrary function of all of the traits, not just a linear summation of an individual payoff function for each trait.

[6] Our notion of multi-dimensional learning should not be confused with Arieli and Mueller-Frank's (2019) different use of the same phrase. In their setup the phrase describes agents who take actions sequentially, and the order in which actions are taken is determined by a multi-dimensional integer lattice rather than a line as in the standard model of herding.

[7] See also Herold and Netzer (2023); Steiner and Stewart (2016); Netzer et al. (2021) on stable compensatory biases.

its advantage in large genotype spaces: while asexual populations sample once and find a local maximum, sexual populations continuously re-sample, improving their chances of reaching a globally superior outcome. Palaiopanos et al. (2017) study the polynomial multiplicative weights update (MWU) algorithm, showing that interior trajectories converge to pure Nash equilibria in congestion games where each player uses MWU. Edhan et al. (2021) extend this, noting that the replicator dynamic is a special case of MWU and that genetic recombination can be viewed as a potential game.

A key difference between our model and the existing literature is that the latter focuses on situations in which a player's payoff depends only on her own action. The strategic aspect of our model, where an agent's payoff depends on others' behaviour, leads to qualitatively different results.

3 Model

Let $G = (A, u)$ be a two-player symmetric normal form game, where A is a (finite) set of actions and $u : A^2 \to \mathbb{R}_{++}$ is a payoff function. We endow actions with structure, positing that each action is a combination of various strategic traits across several dimensions of behaviour. Namely, we posit that there is a set $D = \{1, \ldots, |D|\}$ of dimensions with $|D| \geq 2$ and for every dimension $d \in D$, there is a finite set A_d of *traits* such that $A := \prod_{d \in D} A_d$. Each action $\mathbf{a} \in A$ is thus a D-tuple (vector), i.e., $\mathbf{a} = (a_d)_{d \in D} = (a_1, \ldots, a_{|D|})$. The interpretation is that in each dimension d the traits are mutually exclusive – an agent exhibits a single trait $a_d \in A_d$ in every dimension $d \in D$.

We interpret $u(\mathbf{a}, \mathbf{a}') \in \mathbb{R}_{++}$ as the payoff of an agent playing action $\mathbf{a}$ against one playing action $\mathbf{a}'$. A continuum of agents of mass one is presumed. Each agent is associated with an action $\mathbf{a} \in A$, and the population state (abbr., *state*) x is the distribution of actions in the population, and is an element of the simplex $\Delta(A)$, thus $x(\mathbf{a})$ is the frequency of agents in the population playing $\mathbf{a}$. We demonstrate the multi-dimensionality of actions in the following simple motivating example.

Example 1 (Complementary improvements). There are two dimensions, denoted m and p. Each dimension has two traits, denoted 0 and 1. Thus, the four actions are 0_m0_p, 0_m1_p,

1_m0_p, and 1_m1_p. Payoffs are independent of the opponent's action and are given by

Action	0_m0_p	0_m1_p	1_m0_p	1_m1_p
Payoff	5	1	1	6

The example captures the stylised case in which changing either dimension alone is harmful, whereas changing both dimensions together is beneficial. ♦

We consider a continuous-time setting where incumbents are replaced at a constant rate normalised to 1. Equivalently, this represents a scenario where agents occasionally revise their actions. Let $x_t \in \Delta(A)$ denote the state at time $t \geq 0$. Under the replicator dynamics benchmark, each new agent adopts an action by imitating an existing one, with the likelihood of imitation increasing with the action's success and frequency. Specifically, the probability of choosing action a is proportional to its payoff multiplied by its frequency.

Formally, let $u_x(\mathbf{a}) = \sum_{a' \in A} x(\mathbf{a}')u(\mathbf{a}, \mathbf{a}')$ denote the mean payoff of action $\mathbf{a}$ in state x, and let $u_x = \sum_{\mathbf{a} \in A} x(\mathbf{a})u_x(\mathbf{a})$ denote the mean payoff of the population. The replicator dynamics is defined as follows:[8]

$$(1) \qquad \dot{x}_t(\mathbf{a}) = \frac{x_t(\mathbf{a})u_{x_t}(\mathbf{a})}{u_{x_t}} - x_t(\mathbf{a}) \qquad \textit{(replicator dynamics).}$$

We interpret Equation (1) as follows. Players occasionally revise their actions at a unit rate. When a player revises, they randomly select another agent and imitate their action, with the probability of selecting each agent being proportional to that agent's payoff.

The combinator dynamic changes the object of imitation from multi-dimensional actions to traits. Each new agent evaluates traits independently across dimensions. The probability of imitating a trait in a given dimension increases with its success and frequency. Specifically, the likelihood of selecting trait a_d is proportional to its payoff multiplied by its frequency. After selecting traits for each dimension, the agent combines them

[8] Our functional form slightly differs from the most commonly-used functional form for the replicator dynamics by a payoff-dependent rescaling of time; see Weibull (1997, Section 4.4.3). Weibull also presents a more general imitation dynamics in which $u_x(\mathbf{a})$ (resp., u_x) in Equation (1) is replaced by $w(u_x(\mathbf{a}))$ (resp., $w(u_x)$), where $w : \mathbb{R} \to \mathbb{R}^{++}$ is a strictly monotone function. Our dynamics can capture this general version by a normalisation of the payoff function. That is, if the original payoff function is denoted $\pi : A \to \mathbb{R}$ (which might be measured in dollars), then $u \equiv w(\pi)$ is the normalised payoff following a monotone transformation to cardinal units, measuring the probability of being imitated.

to determine their adopted action. Formally, let $\text{supp}(x) = \{\mathbf{a} \in A | x(\mathbf{a}) > 0\}$ denote the set of actions with a positive probability in state x, and let $\text{supp}_d(x) = \{a_d \in A_d | x(a_d) > 0\}$ denote the set of traits in dimension d with a positive probability in state x. For every trait $a_d \in \text{supp}_d(x)$, let $x(a_d) := \sum_{\mathbf{a}_{-d} \in A_{-d}} x(a_d, \mathbf{a}_{-d})$ denote its (marginal) frequency, and let $u_x(a_d)$ denote its mean (marginal) payoff:

$$u_x(a_d) = \sum_{\mathbf{a}_{-d} \in A_{-d}} \frac{x(a_d, \mathbf{a}_{-d})}{x(a_d)} u_x(a_d, \mathbf{a}_{-d}).$$

The combinator dynamics is defined as follows:

$$\dot{x}_t(\mathbf{a}) = \prod_{d \in D} \frac{x_t(a_d) u_{x_t}(a_d)}{u_{x_t}} - x_t(\mathbf{a}), \qquad \textit{(combinator dynamics)} \tag{2}$$

We interpret Equation (2) as follows. Players occasionally revise their actions at a unit rate. When a player revises, they consider each dimension separately. In each dimension, they independently select another agent at random and imitate that agent's trait in the corresponding dimension, with the probability of selecting each agent being proportional to the agent's payoff.

To simplify notation, we omit the subscript t where it causes no confusion (e.g., writing $\dot{x}$ instead of $\dot{x}_t$). It is straightforward to verify that the combinator dynamics (and similarly, its extension to the recombinator dynamics in Section 5) is an autonomous first-order differential equation and that the vector field defining it is continuous and locally-Lipschitz, hence a solution exists and is unique by the Peano and Picard-Lindelöf theorems.[9]

Our baseline model studies the homogeneous case in which all revising agents use the same revision protocol. Section 5 introduces heterogeneous revision, where some agents copy whole actions and others recombine traits.

[9] The Peano Theorem establishes the global existence of the solution. As the dynamics are locally Lipschitz-continuous, the Picard-Lindelöf theorem establishes uniqueness.

4 Results

4.1 Monotonicity

Dynamics are payoff monotone (see, e.g., Weibull 1997, Definition 4.2) if an action with a higher payoff grows at a higher rate.

Definition 1. A dynamic $\dot{x}$ is *payoff monotone* if $\frac{\dot{x}(\mathbf{a})}{x(\mathbf{a})} > \frac{\dot{x}(\mathbf{a}')}{x(\mathbf{a}')} \iff u_x(\mathbf{a}) > u_x(\mathbf{a}')$ for every state $x \in \Delta(A)$ and for every pair of actions $\mathbf{a}, \mathbf{a}' \in \operatorname{supp}(x)$.

A state is *stationary* if it is a fixed point of the dynamics (i.e., $x \in \Delta(A)$ is stationary if $\dot{x}(\mathbf{a}) = 0$ for every $\mathbf{a} \in A$). It is well-known that every convergent limit state of a trajectory $\lim_{t \to \infty} x_t$ must be a stationary state (Proposition 6.3 of Weibull, 1997). The replicator dynamic satisfies payoff monotonicity, which implies that its stationary states are those that satisfy the property that all incumbent actions have the same payoff.

Fact 1. *(Weibull, 1997, Proposition 5.9)*

1. *The replicator dynamic is payoff monotone.*
2. *A state x is stationary if and only if $u_x(\boldsymbol{a}) = u_x(\boldsymbol{a}')$ for all $\boldsymbol{a}, \boldsymbol{a}' \in \operatorname{supp}(x)$.*

Our next example demonstrates that the combinator dynamics violates payoff monotonicity, and that this can allow strictly dominated actions to be asymptotically stable.

Example 1 (continued). Fix a sufficiently small $\varepsilon \ll 1$. Consider an initial state x that puts weight $1 - \varepsilon$ on action $0_m 0_p$ and weight ε on action $1_m 1_p$. Observe that $0_m 0_p$ is dominated by $1_m 1_p$: $u_x(0_m 0_p) = 5 < u_x(1_m 1_p) = 6$. In what follows, we show that $\frac{\dot{x}(0_m 0_p)}{x(0_m 0_p)} > \frac{\dot{x}(1_m 1_p)}{x(1_m 1_p)}$, which violates payoff monotonicity. Note that the average payoff in the population is $u_x = 5 \cdot (1 - \varepsilon) + 6 \cdot \varepsilon = 5 + \varepsilon$. Substituting this into Equation (2) yields:

$$\dot{x}(0_m 0_p) = \left(\frac{5 - 5\varepsilon}{5 + \varepsilon}\right)^2 - (1 - \varepsilon) = O(\varepsilon) \Rightarrow \frac{\dot{x}(0_m 0_p)}{x(0_m 0_p)} = O(\varepsilon),$$

$$\dot{x}(1_m 1_p) = \left(\frac{6\varepsilon}{5 + \varepsilon}\right)^2 - \varepsilon = O(\varepsilon^2) - \varepsilon \Rightarrow \frac{\dot{x}(1_m 1_p)}{x(1_m 1_p)} = -1 + O(\varepsilon).$$

Figure 1: Phase Portrait in Example 1

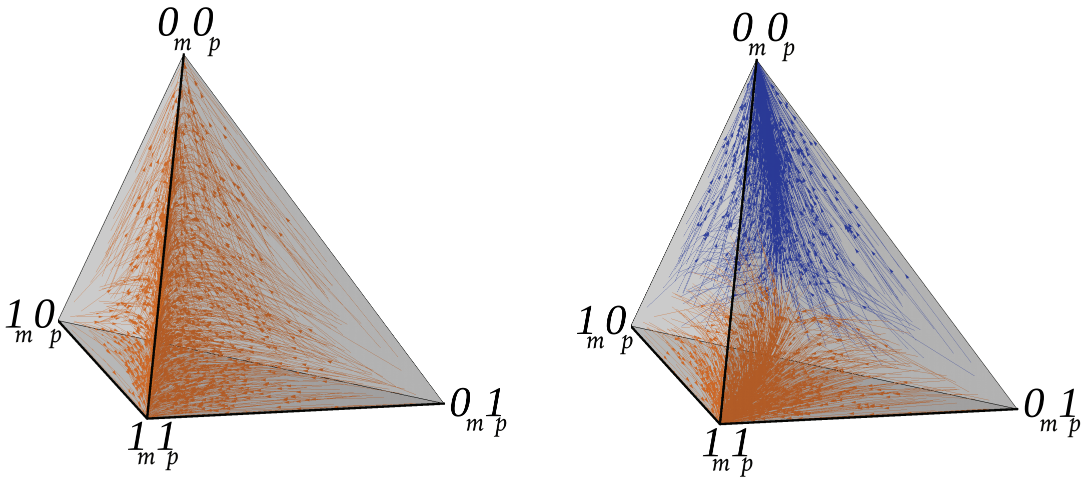

This figure illustrates the phase portraits and the basins of attraction in Example 1 for the replicator dynamics (left panel, where action 1_m1_p is globally stable) and for the combinator dynamics (right panel, where there are two asymptotically stable states, 0_m0_p and 1_m1_p, with their basins of attraction occupying approximately 40% and 60% of the simplex volume, respectively). Blue (resp., orange) arrows represent states that converge to 0_m0_p (resp., 1_m1_p).

Observe that for sufficiently small ε, $\frac{\dot{x}(0_m0_p)}{x(0_m0_p)} > \frac{\dot{x}(1_m1_p)}{x(1_m1_p)}$. Moreover, we later show (Corollary 2) that the environment admits two asymptotically stable states: 0_m0_p and 1_m1_p. Figure 1 illustrates the phase plot and the basins of attraction of 0_m0_p and 1_m1_p. ♦

Although the combinator dynamics do not satisfy the standard payoff monotonicity, they do satisfy the following two types of monotonicity: (1) the action dynamics are monotone with respect to the *product-of-trait utility*

$$\hat{u}_x(\mathbf{a}) := \frac{\prod_{d\in D} x(a_d)}{x(\mathbf{a})} \cdot \prod_{d\in D} u_x(a_d),$$

and (2) the induced trait dynamics are monotone with respect to the standard utility u. Thus, while traits are evaluated by the standard utility, actions are evaluated by a different utility that is the product of two components: (1) the ratio between the product of the traits' frequencies and the action's frequency, and (2) the product of the traits' utilities. Note, in particular, that an action's utility plays no direct role (apart from its indirect impact on the traits' payoffs) in how the action's frequency evolves.

Proposition 1. *The combinator dynamics satisfy:*

1. *The action dynamics are monotone with respect to the product-of-trait utility $\hat{u}$:*

$$\frac{\dot{x}(\boldsymbol{a})}{x(\boldsymbol{a})} > \frac{\dot{x}(\boldsymbol{a}')}{x(\boldsymbol{a}')} \Leftrightarrow \hat{u}_x(\boldsymbol{a}) > \hat{u}_x(\boldsymbol{a}') \Leftrightarrow \frac{\prod_{d\in D} x(a_d)}{x(\boldsymbol{a})} \cdot \prod_{d\in D} u_x(a_d) > \frac{\prod_{d\in D} x(a'_d)}{x(\boldsymbol{a}')} \cdot \prod_{d\in D} u_x(a'_d)$$

for every state $x \in \Delta(A)$ and for every pair of actions $\boldsymbol{a}, \boldsymbol{a}' \in \operatorname{supp}(x)$.

2. *The induced trait dynamics are monotone with respect to the standard utility u:*

$$\frac{\dot{x}(a_d)}{x(a_d)} > \frac{\dot{x}(a'_d)}{x(a'_d)} \iff u_x(a_d) > u_x(a'_d)$$

for every state $x \in \Delta(A)$ and for every pair of traits $a_d, a'_d \in \operatorname{supp}_d(x)$.

Proof. 1. Dividing Eq. (2) by $x(\mathbf{a})$ yields:

$$\frac{\dot{x}(\mathbf{a})}{x(\mathbf{a})} = \frac{1}{x(\mathbf{a})} \prod_{d\in D} \frac{x(a_d)\, u_x(a_d)}{u_x} - 1 = \frac{1}{(u_x)^{|D|}} \frac{\prod_{d\in D} x(a_d)}{x(\mathbf{a})} \cdot \prod_{d\in D} u_x(a_d) - 1,$$

which implies that

$$\frac{\dot{x}(\mathbf{a})}{x(\mathbf{a})} > \frac{\dot{x}(\mathbf{a}')}{x(\mathbf{a}')} \Leftrightarrow \frac{\prod_{d\in D} x(a_d)}{x(\mathbf{a})} \cdot \prod_{d\in D} u_x(a_d) > \frac{\prod_{d\in D} x(a'_d)}{x(\mathbf{a}')} \cdot \prod_{d\in D} u_x(a'_d).$$

2. Slightly rewriting Eq. 2 with $\mathbf{a} = (a_d, \mathbf{a}_{-d})$ yields:

$$\dot{x}(a_d, \mathbf{a}_{-d}) = \frac{x(a_d)\, u_x(a_d)}{u_x^{|D|}} \prod_{d'\neq d} x(a_{d'})\, u_x(a_{d'}) - x(a_d, \mathbf{a}_{-d})$$

Dividing by $x(a_d)$ and summing over all $\mathbf{a}_{-d} \in A_{-d}$ yields:

$$\frac{\dot{x}(a_d)}{x(a_d)} = \sum_{\mathbf{a}_{-d}\in A_{-d}} \frac{\dot{x}(a_d, \mathbf{a}_{-d})}{x(a_d)} = \frac{u_x(a_d)}{u_x^{|D|}} \sum_{\mathbf{a}_{-d}\in A_{-d}} \prod_{d'\neq d} x(a_{d'})\, u_x(a_{d'}) - 1$$

where the sum $\sum_{\mathbf{a}'_{-d}\in A_{-d}} \prod_{d'\neq d} x(a_{d'})\, u_x(a_{d'})$ involves only elements that depend on traits outside dimension d. Hence, it takes the same value for any trait $a'_d \in \operatorname{supp}_d(x)$. This proves part (2). □

4.2 Stationary States

An immediate corollary of Proposition 1 is that a state is stationary iff all actions in its support have the same product-of-trait utility $\hat{u}$:

Corollary 1. State x is stationary under the combinator dynamics if and only if $\hat{u}_x(\mathbf{a}) = \hat{u}_x(\mathbf{a}')$ for all $\mathbf{a}, \mathbf{a}' \in \operatorname{supp}(x)$.

Proof. $\hat{u}_x(\mathbf{a}) = \hat{u}_x(\mathbf{a}')\ \forall \mathbf{a}, \mathbf{a}' \in \operatorname{supp}(x) \iff \frac{\dot{x}(\mathbf{a})}{x(\mathbf{a})} = \frac{\dot{x}(\mathbf{a}')}{x(\mathbf{a}')}\ \forall \mathbf{a}, \mathbf{a}' \in \operatorname{supp}(x) \iff \frac{\dot{x}(\mathbf{a})}{x(\mathbf{a})} = 0\ \forall \mathbf{a} \in \operatorname{supp}(x) \iff \dot{x}(\mathbf{a}) = 0\ \forall \mathbf{a} \in \operatorname{supp}(x) \iff x$ is stationary. □

Next, we show that a state is stationary if and only if (1) all traits obtain the same payoff, and (2) the frequency of each action is the product of the frequencies of its traits.

Proposition 2. *A state x is stationary under the combinator dynamics if and only if*

1. *$u_x(a_d) = u_x$ for any dimension $d \in D$ and any trait $a_d \in \operatorname{supp}_d(x)$, and*

2. *$x(\boldsymbol{a}) = \prod_{d \in D} x(a_d)$ for any action $\boldsymbol{a} \in \operatorname{supp}(x)$.*

Proof. First assume that x is stationary, and we will prove that (1) and (2) hold as follows:

1. It will suffice to prove that $u_x(a_d) = u_x(a'_d)$ for every dimension d and every pair of traits $a_d, a'_d \in \operatorname{supp}_d(x)$, since in such cases $u_x(a_d)$ must equal the mean over traits in $\operatorname{supp}_d(x)$, namely u_x. To demonstrate this, fix d and assume that there is a constant c satisfying $u_x(a_d) = c$ for every $a_d \in \operatorname{supp}_d(x)$. Notice that

$$u_x = \sum_{a_d \in A_d} x(a_d) u_x(a_d) = \sum_{a_d \in \operatorname{supp}_d(x)} x(a_d) c = c,$$

   hence $c = u_x$, as claimed. We next turn to prove the claim. By part (2) of Proposition 1, if, by contradiction, $u_x(a_d) > u_x(a'_d)$ for some dimension d and pair of traits $a_d, a'_d \in \operatorname{supp}_d(x)$, then $\frac{\dot{x}(a_d)}{x(a_d)} > \frac{\dot{x}(a'_d)}{x(a'_d)}$, contradicting the stationarity of x. Thus $u_x(a_d) = u_x$ for every dimension d and trait $a_d \in \operatorname{supp}_d(x)$.

2. As x is stationary and $\mathbf{a} \in \operatorname{supp}(x)$, the combinator Equation (2) becomes

$$0 = \prod_{d \in D} \frac{u_x(a_d)}{u_x} x(a_d) - x(\mathbf{a}).$$

Part (1) of this proposition implies that $\frac{u_x(a_d)}{u_x} = 1$ for every dimension d. Thus the first equation in this part simplifies to

$$0 = \prod_{d \in D} x(a_d) - x(\mathbf{a}),$$

implying property (2).

In the other direction we assume properties (1) and (2) hold and seek to prove that x is stationary. Property (1) implies that $\frac{u_x(a_d)}{u_x} = 1$ for every $\mathbf{a} \in \mathrm{supp}(x)$ and every dimension d, hence the combinator equation, Equation (2), becomes

$$\dot{x}(\mathbf{a}) = \prod_{d \in D} x(a_d) - x(\mathbf{a}) = 0$$

with the last equality to 0 following from property (2). □

Proposition 2 implies that the combinator dynamics are trait-centred (in a way analogous to the gene-centred view of genetic evolution (Williams, 1966; Dawkins, 1976)). The social learning process that is captured by the combinator dynamics leads to the survival of payoff-maximising traits (rather than payoff-maximising actions), where each trait $a_d \in A_d$ is essentially competing against the other traits in A_d. Moreover, the frequency of each action is just the product of the frequencies of its traits.

Observe that actions in a stationary state may have different payoffs. This is illustrated in the following two examples.

Example 1 (continued). Consider the following full-support state: $\hat{x}(0_m 0_p) = \left(\frac{5}{9}\right)^2$, $\hat{x}(1_m 1_p) = \left(\frac{4}{9}\right)^2$, $\hat{x}(1_m 0_p) = \hat{x}(0_m 1_p) = \frac{5}{9} \cdot \frac{4}{9}$. Observe that the frequencies of the traits in each dimension are: $\hat{x}(0_m) = \hat{x}(0_p) = \frac{5}{9}$, $\hat{x}(1_m) = \hat{x}(1_p) = \frac{4}{9}$, and that the frequency of each action is the product of the traits' frequencies. Further note that all the traits have the same payoff: $u_{\hat{x}}(0_m) = u_{\hat{x}}(0_p) = \frac{4}{9} \cdot 1 + \frac{5}{9} \cdot 5 = \frac{29}{9}$, $u_{\hat{x}}(1_m) = u_{\hat{x}}(1_p) = \frac{5}{9} \cdot 1 + \frac{4}{9} \cdot 6 = \frac{29}{9}$. This implies that state $\hat{x}$ is stationary (even though different actions obtain different payoffs). ♦

Example 2 (Congestion with cross-dimensional fit). Similarly to Example 1, we again have two dimensions, denoted by m and p, and in each dimension there are two traits, denoted by 0 and 1. The payoff of an action is the sum of three components:

Table 1: Payoff Matrix for Example 2 (Congestion with Cross-Dimensional Fit)

	$0_m 0_p$	$0_m 1_p$	$1_m 0_p$	$1_m 1_p$
$0_m 0_p$	3^{3}	8^{6}	8^{6}	13^{13}
$0_m 1_p$	6^{8}	1^{1}	11^{11}	6^{8}
$1_m 0_p$	6^{8}	11^{11}	1^{1}	6^{8}
$1_m 1_p$	13^{13}	8^{8}	8^{6}	3^{3}

1. *dimension-m congestion*: five times the share of agents choosing the other trait in dimension m; for example, choosing 1_m contributes $5 \cdot x(0_m) = 5 \cdot (1 - x(1_m))$.

2. *dimension-p congestion*: five times the share of agents choosing the other trait in dimension p; for example, choosing 1_p contributes $5 \cdot x(0_p) = 5 \cdot (1 - x(1_p))$.

3. *cross-dimensional fit*: a payoff of 3 is added when the two traits are aligned, $1_m 1_p$ or $0_m 0_p$, and a payoff of 1 is added when they are not aligned, $1_m 0_p$ or $0_m 1_p$.

The example captures a stylised situation in which traits are valuable when they are scarce in their own dimension, but full actions also require fit across dimensions. Thus, a trait may look locally attractive even when the full bundle in which it appears is relatively low-performing. The payoff matrix is illustrated in Table 1. Specifically,

$$u_x(1_m 1_p) = 5 \cdot (x(0_m) + x(0_p)) + 3, \quad u_x(0_m 0_p) = 5 \cdot (x(1_m) + x(1_p)) + 3,$$
$$u_x(1_m 0_p) = 5 \cdot (x(0_m) + x(1_p)) + 1, \quad u_x(0_m 1_p) = 5 \cdot (x(1_m) + x(0_p)) + 1.$$

The replicator dynamic admits a unique non-pure stationary state x^*, in which half of the population uses $1_m 1_p$ and the other half uses $0_m 0_p$. Proposition 2 implies that the combinator dynamic admits a unique non-pure stationary state $\hat{x}$, in which each of the four actions has frequency 0.25. Both states are globally stable, as illustrated in Figure 2.

At $\hat{x}$, the aligned actions $0_m 0_p$ and $1_m 1_p$ have payoff 8, while the non-aligned actions $0_m 1_p$ and $1_m 0_p$ have payoff 6. Nevertheless, all four traits have the same payoff, equal to 7. Thus, component-wise learning does not eliminate the low-payoff actions: each trait

Figure 2: Phase Portrait in Example 2.

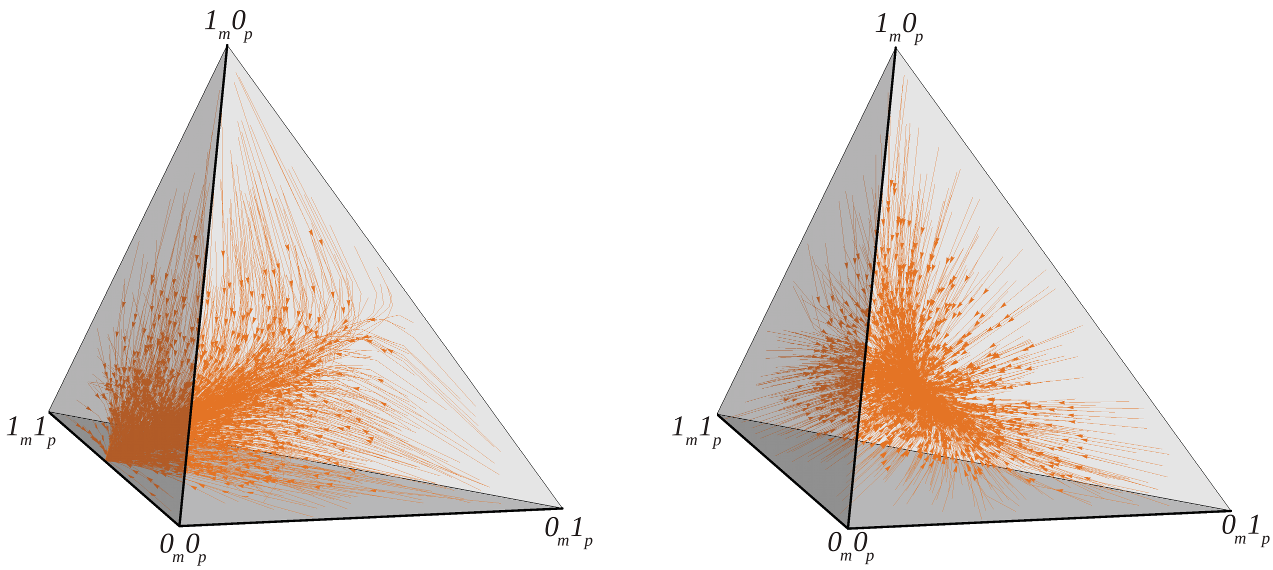

This figure illustrates the phase portraits in Example 2 for the replicator dynamics (left panel, where state $x^* = 0.5 \cdot 1_m 1_p + 0.5 \cdot 0_m 0_p$ is globally stable) and for the combinator dynamics (right panel, where the uniform state $\hat{x}$ is globally stable).

appears just as successful as its alternative when evaluated in its own dimension. ◆

4.3 Stable States

A state is Lyapunov stable if a population trajectory starting near it always remains close.

Definition 2. A state $x^* \in \Delta(A)$ is *Lyapunov stable* if for every neighbourhood U of x^* there is a neighbourhood $V \subseteq U$ of x^* such that if the initial state of a trajectory satisfies $x_0 \in V$ then $x_t \in U$ for all $t > 0$. A state is *unstable* if it is not Lyapunov stable.

As shown in Weibull, 1997, Section 6.4, every Lyapunov stable state is stationary. A Lyapunov stable state is asymptotically stable if starting from every sufficiently close initial condition, the trajectories converge to that state.

Definition 3. A state $x^* \in \Delta(A)$ is *asymptotically stable* if (1) it is Lyapunov stable and (2) there is an open neighbourhood U of x^* such that all trajectories initially in U converge to x^*, i.e., $x_0 \in U \Rightarrow \lim_{t\to\infty} x_t = x^*$.

Finally, we say that state x^* is *globally stable* if any trajectory initially in an interior state converges to x^*, i.e., $x^0 \in \mathrm{Int}(\Delta(A)) \Rightarrow \lim_{t\to\infty} x^t = x^*$. It is well known that asymptotic stability of stationary states under the replicator dynamic is characterised by two conditions:

1. *Internal stability*: the payoff matrix restricted to the incumbent actions is (semi-) negative-definite with respect to the tangent space.

2. *External stability*: The payoffs of external actions are lower than those of incumbents.

The strict variants of these conditions (i.e., negative-definiteness and quasi-strictness) imply asymptotic stability, and their weak counterparts (semi-negative-definiteness and being a Nash equilibrium) are implied by asymptotic stability. This is formalised as follows. Let J^u be the *Jacobian matrix* of the payoff function; that is, J^u is a square matrix of size $|\operatorname{supp}(x)|$, where the ij-th element in the matrix is

$$(J^u)_{\mathbf{a}^i,\mathbf{a}^j} = \frac{\partial u_x(a^i)}{\partial x(a^j)} = u(\mathbf{a}^i,\mathbf{a}^j) \qquad \text{for all } \mathbf{a}^i,\mathbf{a}^j \in \operatorname{supp}(x),$$

where we arbitrarily order the actions in A as $(a^1, a^2, ..., a^{|A|})$.

Fact 2. 1. *If a state x is asymptotically stable under the replicator dynamics, then it satisfies*

   (a) weak internal stability: $\boldsymbol{w}^\intercal \cdot J^u \cdot \boldsymbol{w} \leq 0$ for every $\boldsymbol{w} \in T_{\operatorname{supp}(x)}$; and

   (b) weak external stability: $u_x \geq u_x(\boldsymbol{a})$ for every $\boldsymbol{a} \notin \operatorname{supp}(x)$.

2. *A stationary state x is asymptotically stable under the replicator dynamics if the strict counterparts of conditions (a) and (b) are satisfied.*

We omit the proof of this well-known result (which is implied by combining Theorem 9.2.7, Corollary 9.4.2, Theorem 9.4.4, and Theorem 9.4.8 in Van Damme, 1991).

The results of this subsection characterise asymptotic stability under the combinator dynamics. In order to do so, we first need to define the $\hat{u}$-Jacobian matrix and the payoff of an external trait under the combinator dynamics.

For every state x, let $J_x^{\hat{u}}$ denote the Jacobian of the product-of-trait utility $\hat{u}$ at x (henceforth, the *$\hat{u}$-Jacobian matrix*); that is, $J_x^{\hat{u}}$ is a square matrix of size $|\operatorname{supp}(x)|$, where the ij-th element in the matrix is the partial derivative of $\hat{u}_x(\mathbf{a}_i)$ with respect to $x(\mathbf{a}_j)$:

$$(3) \qquad (J_x^{\hat{u}})_{\mathbf{a}^i,\mathbf{a}^j} = \frac{\partial \hat{u}_x(\mathbf{a}^i)}{\partial x(\mathbf{a}^j)} \qquad \text{for all } a^i, a^j \in \text{supp}(x).$$

For every state x and every external trait $a_d \notin \text{supp}_d(x)$, let $u_x(a_d)$ be the payoff of trait a_d when interacting with traits in other dimensions, drawn according to x. Formally, define the marginal frequency of $\mathbf{a}_{-d}$ as $x(\mathbf{a}_{-d}) = \sum_{a'_d \in A_d} x(a'_d, \mathbf{a}_{-d})$. Then $u_x(a_d)$ is:

$$u_x(a_d) = \sum_{\mathbf{a}_{-d} \in A_{-d}} x(\mathbf{a}_{-d}) \cdot u_x(a_d, \mathbf{a}_{-d}).$$

In what follows we characterise the set of asymptotically stable states under the combinator dynamics in terms of two conditions analogous to Fact 2:

1. *Internal stability:* The $\hat{u}$-Jacobian matrix is (semi-)negative-definite.

2. *External stability:* The payoffs of external traits are lower than those of incumbents.

Our main result is, essentially, an if and only if characterisation of asymptotic stability. Specifically, the strict variants of these conditions imply asymptotic stability, and their weak counterparts are implied by Lyapunov stability.

Theorem 1. *1. If a state x^* is Lyapunov stable under the combinator dynamics, then it satisfies*

   (a) Weak internal stability: $\boldsymbol{w}^\intercal \cdot J_{x^*}^{\hat{u}} \cdot \boldsymbol{w} \leq 0$ for every $\boldsymbol{w} \in T_{\text{supp}(x^*)}$; and

   (b) Weak external stability against traits: $u_{x^*} \geq u_{x^*}(a_d)$ for every $a_d \notin \text{supp}_d(x^*)$.

2. A stationary state x^* is asymptotically stable under the combinator dynamics if the strict counterparts of conditions (a) and (b) hold.

Theorem 1 is a special case of Theorem 1′ that is formally proven in Appendix A.3. We henceforth present the sketch of proof.

Sketch of proof for Theorem 1. (Special case of Theorem 1′).

1. (a) Suppose, for contradiction, that the $\hat{u}$-Jacobian matrix $J_{x^*}^{\hat{u}}$ is not negative semi-definite. Then, there exists a state y with $\text{supp}\, y \subset \text{supp}\, x^*$, such that $(x^* - y)' J_{x^*}^{\hat{u}} (x^* -$

$y) > 0$. For small $\varepsilon > 0$, the perturbed state $\varepsilon y + (1-\varepsilon)x^*$ yields a higher weighted average product-of-trait utility $\hat{u}$ under y than under x^*. By monotonicity in $\hat{u}$ (part (i) of Proposition 1), the population moves away from x^*, a contradiction.

(b) Suppose, for contradiction, that $u_{x^*} < u_{x^*}(\hat{a}_d)$ for some $\hat{a}_d \notin \text{supp}_d(x^*)$. Define y as a distribution assigning positive mass only to actions with trait $\hat{a}_d$ in dimension d. For such actions $\mathbf{a}$, we define $y(\mathbf{a}) = x^*(\mathbf{a}_{-d})$. In the perturbed state $\varepsilon y + (1-\varepsilon)x^*$ for small $\varepsilon > 0$, trait $\hat{a}_d$ has payoff $u_{x^*}(\hat{a}_d)$ that is strictly higher than u_{x^*}. Due to trait monotonicity, this implies that the frequency of $\hat{a}_d$ increases, and that the population moves away from x^*, a contradiction.

2. Assume to the contrary that x^* is not asymptotically stable. Then for every $\varepsilon > 0$ there is an ε-nearby state y_ε (i.e., satisfying $|y_\varepsilon - x^*| < \varepsilon$) such that a population starting at y_ε does not converge to x^*. The fact that the $\hat{u}-$Jacobian matrix is negative definite implies that $y_\varepsilon(\mathbf{a}) \to x^*(\mathbf{a})$ for every $a \in \text{supp}(x^*)$ as $\varepsilon \to 0$. The inequality $u_{x^*} > u_{x^*}(a_d)$ for every $a_d \notin \text{supp}_d(x^*)$, implies, due to trait monotonicity (Proposition 1) that the share of agents who play each external trait a_d converges to zero. Combining these facts, we see that the population converges from any sufficiently nearby state y_ε to x^*, which implies that x^* is asymptotically stable. □

Next, we apply Theorem 1 to characterise stable states in two interesting special cases: (I) pure states, and (II) interior states. Corollary 2 shows that a pure state $\mathbf{a}$ is stable iff the payoff of $\mathbf{a}$ against itself is higher than the payoff of any "neighbouring" action (which differs in a single trait) against $\mathbf{a}$.

Corollary 2.
1. If a pure state $\mathbf{a} \in A$ is Lyapunov stable under the combinator dynamics, then $u_a(\mathbf{a}) \geq u_a(a'_d, \mathbf{a}_{-d})$ for any dimension $d \in D$ and any trait $a'_d \in A_d$.

2. A pure state $\mathbf{a} \in A$ is asymptotically stable under the combinator dynamics if $u_a(\mathbf{a}) > u_a(a'_d, \mathbf{a}_{-d})$ for any dimension $d \in D$ and any trait $a_d \neq a'_d \in A_d$.

Corollary 2 implies that, in Example 1, both 0_m0_p and 1_m1_p are asymptotically stable (while 1_m0_p and 0_m1_p are unstable), and that all pure states in Example 2 are unstable.

Corollary 3 shows that an interior (full support) state x is asymptotically stable essentially if and only if the $\hat{u}$-Jacobian matrix is negative definite.

Corollary 3. 1. If an interior state x^* is Lyapunov stable under the combinator dynamics, then $\mathbf{w}^T J^{\hat{u}}_{x^*} w \leq 0$ for every $\mathbf{w} \in T_A$.

2. An interior stationary state $x^* \in \Delta(A)$ is asymptotically stable under the combinator dynamics if $\mathbf{w}^T J^{\hat{u}}_{x^*} w < 0$ for every $\mathbf{w} \in T_A$.

Corollaries 2–3 follow immediately from Theorem 1. A simple numerical calculation shows that the $\hat{u}$-Jacobian matrices of the unique interior stationary states in Examples 1 & 2 are (where we arbitrarily order the actions as $(0_m 0_p, 1_m 0_p, 0_m 1_p, 1_m 1_p)$):

$$\begin{pmatrix} -0.76 & -0.06 & -0.06 & -3.72 \\ -0.31 & -3.41 & -0.62 & 0.46 \\ -0.31 & -0.62 & -3.41 & 0.46 \\ -3.1 & 0.08 & 0.08 & -0.41 \end{pmatrix} \qquad \begin{pmatrix} -2.43 & 0 & 0 & -1.58 \\ 0 & -3 & -1 & 0 \\ 0 & -1 & -3 & 0 \\ -1.58 & 0 & 0 & -2.43 \end{pmatrix}$$

Jacobian in Example 1 Jacobian in Example 2

The Jacobian in Example 1 (complementary improvements) is not negative definite with respect to the tangent space (in particular, $\mathbf{w}^T J^{\hat{u}}_{x^*} \mathbf{w} = 5.65 > 0$ for $\mathbf{w}^T = (1, 0, 0, -1)$). By contrast, the Jacobian in Example 2 (congestion-with-fit game) is negative-definite with respect to the tangent space, which implies that $\hat{x}$ is asymptotically stable. Thus, a stable state may contain persistent payoff differences across actions, even though all incumbent traits have equal payoff.

5 Recombinator Dynamics (Heterogeneous Environments)

In this section, we extend our model by allowing settings where some agents imitate whole actions, while others combine imitated traits. Specifically, we assume that a share of $r \in [0,1]$ of the new agents combine imitated traits, while the remaining share of $1 - r$ replicate entire actions. We call this parameter r the *recombination rate*. Formally, the *recombinator dynamics* with parameter r are given by

$$\text{(4)} \qquad \dot{x}(\mathbf{a}) = (1 - r)\frac{x(\mathbf{a}) u_x(\mathbf{a})}{u_x} + r \prod_{d \in D} \frac{x(a_d) u_x(a_d)}{u_x} - x(\mathbf{a}) \qquad (\textit{recombinator dynamics}).$$

Observe that $r = 0$ coincides with the replicator dynamics, while $r = 1$ coincides with the combinator dynamics.

5.1 Monotonicity

We define the *r-utility* in state x as follows:

$$u_x^r(\mathbf{a}) := (1-r) \cdot \frac{u_x(\mathbf{a})}{u_x} + r \cdot \frac{\prod_{d \in D} x(a_d)}{x(\mathbf{a})} \cdot \prod_{d \in D} \frac{u_x(a_d)}{u_x}. \tag{5}$$

Next we extend Fact 1 and Proposition 1 and show that under recombinator dynamics for all values of r: (1) the action dynamics are monotone with respect to the r-utility, and (2) the induced trait dynamics are monotone with respect to the standard utility u.

Proposition 1′. *For $r \in (0,1)$, under the recombinator dynamics:*

1. *The action dynamics are monotone with respect to the r-utility u_x^r:*

$$\frac{\dot{x}(\boldsymbol{a})}{x(\boldsymbol{a})} > \frac{\dot{x}(\boldsymbol{a}')}{x(\boldsymbol{a}')} \Leftrightarrow u_x^r(\boldsymbol{a}) > u_x^r(\boldsymbol{a}')$$

   for every state $x \in \Delta(A)$ and for every pair of actions $\boldsymbol{a}, \boldsymbol{a}' \in \operatorname{supp}(x)$.

2. *The induced trait dynamics are monotone with respect to the standard utility u:*

$$\frac{\dot{x}(a_d)}{x(a_d)} > \frac{\dot{x}(a_d')}{x(a_d')} \Leftrightarrow u_x(a_d) > u_x(a_d')$$

   for every state $x \in \Delta(A)$ and for every pair of traits $a_d, a_d' \in \operatorname{supp}_d(x)$.

The proof idea is similar to Proposition 1. The formal proof is presented in Appendix A.1.

5.2 Stationary States

An immediate corollary of Proposition 1′ is that a state is stationary iff all actions in its support have the same r-utility u_x^r:

Corollary 1′. State x is stationary under the recombinator dynamics with parameter $r \in [0,1]$ if and only if $u_x^r(\mathbf{a}) = u_x^r(\mathbf{a}')$ for all $\mathbf{a}, \mathbf{a}' \in \operatorname{supp}(x)$.

Proof. $u_x^r(\mathbf{a}) = u_x^r(\mathbf{a}')\ \forall \mathbf{a}, \mathbf{a}' \in \operatorname{supp}(x) \iff \frac{\dot{x}(\mathbf{a})}{x(\mathbf{a})} = \frac{\dot{x}(\mathbf{a}')}{x(\mathbf{a}')}\ \forall \mathbf{a}, \mathbf{a}' \in \operatorname{supp}(x) \iff \frac{\dot{x}(\mathbf{a})}{x(\mathbf{a})} = 0\ \forall \mathbf{a} \in \operatorname{supp}(x) \iff \dot{x}(\mathbf{a}) = 0\ \forall \mathbf{a} \in \operatorname{supp}(x) \iff x$ is stationary. □

Next, we show that a state is stationary if and only if (1) all traits obtain the same payoff, and (2) all actions obtain the same r-weighted average of the relative action payoff (i.e., the ratio $\frac{u_x(\mathbf{a})}{u_x}$ between the action's payoff and the average payoff) and the trait-to-action frequency ratio ($\frac{\prod_{d \in D} x(a_d)}{x(\mathbf{a})}$).

Proposition 2′. *A state x is stationary under the recombinator dynamics with parameter r iff*

1. *$u_x(a_d) = u_x$ for any dimension $d \in D$ and any trait $a_d \in \operatorname{supp}_d(x)$, and*
2. *Balance of payoffs and frequencies: $(1-r)\frac{u_x(\boldsymbol{a})}{u_x} + r\frac{\prod_{d \in D} x(a_d)}{x(\boldsymbol{a})} = 1$ for any $\boldsymbol{a} \in \operatorname{supp}(x)$.*

The argument is similar to Proposition 2. The formal proof is presented in Appendix A.2.

Example 2 (continued). We apply the recombinator dynamics with parameter r to Example 2. The symmetries between the two dimensions and between the two traits imply that the unique stationary non-pure state x^r is of the form $x^r(0_m 0_p) = x^r(1_m 1_p) = 0.25 + q$ and $x^r(0_m 1_p) = x^r(1_m 0_p) = 0.25 - q$. Observe that: $x^r(0_m) = x^r(1_m) = x^r(0_p) = x^r(1_p) = 0.5$, all traits obtain the same payoff of 7 (i.e., $u_{x^r}(0_m) = u_{x^r}(0_p) = u_{x^r}(1_m) = u_{x^r}(1_p) = 7$), and the average payoff in the population is $u_{x^r} = 2 \cdot (0.25 + q) \cdot 8 + 2 \cdot (0.25 - q) \cdot 6 = 7 + 4q$. This implies that x^r satisfies condition 1 in Proposition 2′. Condition 2 of Proposition 2′ is satisfied if and only if

$$1 = (1-r)\frac{u_{x^r}(0_m 0_p)}{u_{x^r}} + r\frac{x^r(0_m) \cdot x^r(0_p)}{x^r(0_m 0_p)} = \frac{(1-r)\,8}{7+4q} + \frac{0.25r}{0.25+q} \Leftrightarrow q = \frac{\sqrt{49r^2 - 4r + 4} - 7r}{8}.$$

Note that q is decreasing in r, and that $q(r=0) = 0.25$ and $q(r=1) = 0$. The average payoff in the population $u_{x^r} = 7 + 4q$ is decreasing in r. The aligned actions $0_m 0_p$ and $1_m 1_p$ obtain a high payoff of 8, which is balanced by their low trait-to-action frequency ratio of $\frac{0.25}{0.25+q}$. The remaining actions $0_m 1_p$ and $1_m 0_p$ obtain a low payoff of 6, which is balanced by their high trait-to-action frequency ratio of $\frac{0.25}{0.25-q}$. Thus, whole-action imitation favours coherent high-payoff bundles, whereas component-wise imitation pushes the state towards independent trait frequencies. ♦

5.3 Stable States

Next, we characterise the asymptotically stable states under the recombinator dynamics. The characterisation first requires us to define the r-Jacobian for internal stability and the payoff of external traits u_x^r for external stability.

For every state x, let J_x^r denote the Jacobian of the r-utility u_x^r at x (henceforth, the *r-Jacobian matrix*); that is, J_x^r is a square matrix of size $|\operatorname{supp}(x)|$, where the ij-th element in the matrix is the partial derivative of $u_x^r(\mathbf{a}_i)$ with respect to $x(\mathbf{a}_j)$:

$$(J_x^r)_{\mathbf{a}^i,\mathbf{a}^j} = \frac{\partial u_x^r(\mathbf{a}^i)}{\partial x(\mathbf{a}^j)} \qquad \text{for all } \mathbf{a}^i, \mathbf{a}^j \in \operatorname{supp}(x). \tag{6}$$

Consider an invasion of state x by mutants bearing trait $a_d \notin \operatorname{supp}_d(x)$. Under the combinator dynamics ($r = 1$), the invading trait a_d interacts with the incumbent traits in other dimensions according to the distribution x. By contrast, when $r \in (0,1)$, the distribution of trait profiles in other dimensions that interact with the invading trait (henceforth, partners of the invading trait) converges towards a unique stable distribution of partners $\eta_x^{a_d} \neq x$. Let $\operatorname{supp}_{-d}(x)$ denote the trait profiles in all other dimensions except d that are in the support of the population, i.e.,

$$\operatorname{supp}_{-d}(x) = \{\mathbf{a}_{-d} \in A_{-d} \mid \exists a'_d,\ s.t.\ x(a'_d, \mathbf{a}_{-d}) > 0\}.$$

The distribution $\eta_x^{a_d}$ induces each type $(a_d, \mathbf{a}_{-d})$ (where $\mathbf{a}_{-d} \in \operatorname{supp}_{-d}(x)$) with a frequency that is an r-weighted average of two elements: (1) its own payoff times its own frequency, and (2) the product of the frequencies of traits in $\mathbf{a}_{-d}$.

Definition 4. Fix $r > 0$, a stationary state x, and a trait $a_d \notin \operatorname{supp}_d(x)$. Then $\eta_x^{a_d} \in \Delta(\operatorname{supp}_{-d}(x))$, which we call the *stable partner distribution* of trait a_d, is the unique solution to the following set of $|\operatorname{supp}_{-d}(x)|$ equations, one for every $\mathbf{a}_{-d} \in \operatorname{supp}_{-d}(x)$:

$$\eta(\mathbf{a}_{-d}) = (1-r)\frac{\eta(\mathbf{a}_{-d})u_x(a_d, \mathbf{a}_{-d})}{\sum_{\mathbf{a}'_{-d} \in \operatorname{supp}_{-d}(x)} \eta(\mathbf{a}'_{-d})u_x(a_d, \mathbf{a}'_{-d})} + r\prod_{d' \neq d} x(a_{d'}). \tag{7}$$

We begin by showing that the stable partner distribution is well-defined.[10]

Proposition 3. *Equation (7) admits a unique solution in* $\Delta(\mathrm{supp}_{-d}(x))$ *for any* $r > 0$*, stationary state* x*, and trait* $a_d \notin \mathrm{supp}_d(x)$*.*

Proof. The result is immediate when $r = 1$ (in which case $\eta(\mathbf{a}_{-d}) = \prod_{d' \neq d} x(a_{d'})$). We henceforth assume that $r < 1$. A solution for Equation (7) clearly exists by Brouwer's Fixed Point Theorem. To prove uniqueness, we will show that, for a given r, x, and a_d, there is a unique $z_0 > 0$, such that every solution η of Equation (7) satisfies

$$\sum_{\mathbf{a}_{-d} \in \mathrm{supp}_{-d}(x)} \eta(\mathbf{a}_{-d}) u_x(a_d, \mathbf{a}_{-d}) = z_0.$$

Assuming that this holds true, substituting z_0 back into Equation (7) we obtain

$$\eta(\mathbf{a}_{-d}) = (1-r)\frac{\eta(\mathbf{a}_{-d}) u_x(a_d, \mathbf{a}_{-d})}{z_0} + r\prod_{d' \neq d} x(a_{d'}),$$

for every $\mathbf{a}_{-d} \in \mathrm{supp}_{-d}(x)$, which clearly has a unique solution. As this must hold for every solution η, uniqueness follows. To prove the existence of such a z_0, for every $\mathbf{a}_{-d}$ and each $z > (1-r)u_x(a_d, \mathbf{a}_{-d})$, define $\eta_z(\mathbf{a}_{-d})$ by:

$$\eta_z(\mathbf{a}_{-d}) = \frac{z \cdot r \cdot \prod_{d' \neq d} x(a_{d'})}{z - (1-r)u_x(a_d, \mathbf{a}_{-d})}. \tag{8}$$

Let $\bar{z} = \max_{\mathbf{a}_{-d} \in \mathrm{supp}_{-d}(x)} \left((1-r)u_x(a_d, \mathbf{a}_{-d})\right)$. Observe that $\eta_z(\mathbf{a}_{-d})$ is well-defined for every $\mathbf{a}_{-d} \in \mathrm{supp}_{-d}(x)$ and each $z > \bar{z}$. Further observe that $\eta_z(\mathbf{a}_{-d})$ is decreasing in z for every $\mathbf{a}_{-d}$, hence the sum $h(z) = \sum_{\mathbf{a}_{-d} \in \mathrm{supp}_{-d}(x)} \eta_z(\mathbf{a}_{-d})$ is decreasing. Letting $z \to \bar{z}$ from above gives $h(z) \to \infty$ in the limit, and in contrast as $z \to \infty$ one obtains

$$h(z) \to r \sum_{\mathbf{a}_{-d} \in \mathrm{supp}_{-d}(x)} \prod_{d' \neq d} x(a_{d'}) = r < 1.$$

It follows that there is a unique $z_0 > \bar{z} > 0$ such that $h(z_0) = 1$. Given a solution η of

[10] Appendix A.3.5 gives a method to approximate numerically the stable partner distribution. We show that the stable partner distribution is the globally asymptotically stable equilibrium of some dynamics. Thus, one could compute it by simulating these dynamics.

Equation (7), if we set $z = \sum_{\mathbf{a}_{-d} \in \mathrm{supp}_{-d}(x)} \eta(\mathbf{a}_{-d}) u_x(a_d, \mathbf{a}_{-d})$ then solving the equation for η we obtain $\eta = \eta_z$. Furthermore, $z = z_0$ by the uniqueness of z_0. □

The fact that the partner distribution of an invading trait converges to a stable distribution (as shown in the proof of Theorem 1′) allows us to define the payoff of an invading trait as the weighted average of the actions carrying this mutant trait, where the weights are distributed according to the stable distribution of partners.

Definition 5. Fix a stationary state x and recombination rate r. The *payoff of an invading trait* $a_d \notin \mathrm{supp}_d(x)$ is defined as:

$$u_x^r(a_d) = \sum_{\mathbf{a}_{-d} \in \mathrm{supp}_{-d}(x)} \eta_x^{a_d}(\mathbf{a}_{-d}) \cdot u_x(a_d, \mathbf{a}_{-d}) . \tag{9}$$

The stable partner distribution is degenerate if the stationary state is homogeneous (pure). The following example demonstrates the stable partner distribution for a heterogeneous stationary state in a setup with specialised and versatile traits.

Example 3 (Specialised and versatile traits). There are two dimensions. In the first dimension (m), a player chooses a position, either 0_m or 1_m. In the second dimension (p), the player chooses one of three traits: a trait specialised for 1_m, denoted 1_p; a trait specialised for 0_m, denoted 0_p; or a versatile trait, denoted v_p. The payoff is the sum of two components:

1. *dimension-m congestion*: ten times the share of players choosing the other trait in dimension m; for example, choosing 1_m contributes $10 \cdot x(0_m) = 10 \cdot (1 - x(1_m))$.

2. *cross-dimensional fit*: a payoff of 5 is added when a specialised trait in dimension p is paired with its matching trait in dimension m, $1_m 1_p$ or $0_m 0_p$; a payoff of 1 is added when a specialised trait is mismatched, $1_m 0_p$ or $0_m 1_p$; and a payoff of 4 is added when the versatile trait v_p is used.

The example captures a stylised trade-off between specialisation and versatility. A specialised trait performs very well in one coherent configuration but poorly when recom-

Table 2: Payoff Matrix for Example 3 (Specialised and Versatile Traits)

	0_m0_p	0_m1_p	0_mv_p	1_m0_p	1_m1_p	1_mv_p
0_m0_p	5, 5	5, 1	5, 4	15, 11	15, 15	15, 14
0_m1_p	1, 5	1, 1	1, 4	11, 11	11, 15	11, 14
0_mv_p	4, 5	11, 11	4, 4	14, 11	14, 15	14, 14
1_m0_p	11, 15	11, 11	11, 14	1, 1	1, 5	1, 4
1_m1_p	15, 15	15, 11	15, 14	5, 1	5, 5	5, 4
1_mv_p	14, 15	14, 11	14, 14	4, 1	4, 5	4, 4

bined with the wrong partner trait. A versatile trait performs moderately well across configurations.

The payoff matrix is presented in Table 2. Consider the stationary state x^* that assigns equal mass of 0.5 to 1_mv_p and 0_mv_p. Consider an invasion of this population by a small mass of agents carrying the trait 0_p. Substituting the example's parameters in Equation (7) yields the following partner distribution for the invading trait 0_p as a function of r:

$$\eta(0_m) = (1-r)\frac{10\eta(0_m)}{10\eta(0_m)+6(1-\eta(0_m))} + \frac{r}{2} \Rightarrow \eta_{x^*}^{0_p}(0_m)(r) = \frac{\sqrt{4r^2-r+1}-2r+1}{2}. \tag{10}$$

Observe that $\eta_{x^*}^{0_p}(0_m)(r)$ is decreasing in r, $\eta_{x^*}^{0_p}(0_m)(0)=1$, $\eta_{x^*}^{0_p}(0_m)(\frac{1}{4})=\frac{3}{4}$, and $\eta_{x^*}^{0_p}(0_m)(1)=0.5$. The payoff of the invading trait 0_p is

$$u_{x^*}^r(0_p) = \eta_{x^*}^{0_p}(r)(0_m)\cdot u_{x^*}(0_m0_p) + \eta_{x^*}^{0_p}(r)(1_m)\cdot u_{x^*}(1_m0_p) = \eta_{x^*}^{0_p}(r)(0_m)\cdot 10 + \eta_{x^*}^{0_p}(r)(1_m)\cdot 6,$$

which is larger than incumbents' payoff $u_{x^*}=9$ if and only if $\eta_{x^*}^{0_p}(r)(0_m) > \frac{3}{4} \iff r<\frac{1}{4}$. Hence the stationary state x^* is stable against an invasion of trait 0_p if $r>\frac{1}{4}$ and unstable if $r<\frac{1}{4}$. Moreover, numerical calculations show that for $r>\frac{1}{4}$, state x^* is globally stable, i.e., the population converges to everyone having the versatile trait and receiving the same payoff of 9, starting from any interior state. By contrast, when $0<r<\frac{1}{4}$

the numerical calculations show that the environment admits a globally stable state $\hat{x}^r$ in which the versatile trait is extinct and incumbents with a high payoff of 10 coexist with those receiving a low payoff of 6. (Specifically, $\hat{x}^r(0_m 0_p) = \hat{x}^r(1_m 1_p) = 0.25 + q$ and $\hat{x}^r(0_m 1_p) = \hat{x}^r(1_m 0_p) = 0.25 - q$, where $q(r) = \frac{\sqrt{4 \cdot r^2 - r + 1} - 2r}{4}$.) ♦

Next we extend Fact 2 and Theorem 1, and characterise the set of asymptotically stable states as stationary states that satisfy the following three conditions.

1. *Internal stability:* The r-Jacobian matrix is (semi-)negative-definite.
2. *External stability against traits:* The payoffs of traits outside the support are lower than the payoffs of the traits of the incumbents.
3. *External stability against actions:* The payoffs of actions outside the support are at most $\frac{1}{1-r}$ times higher than the average payoffs of the incumbents.

Our result is, essentially, an if and only if characterisation of asymptotic stability. Specifically, the strict variants of these conditions (i.e., negative-definiteness and strictly lower payoffs of external traits/actions) imply asymptotic stability, and their weak counterparts (semi-negative-definiteness and weakly lower payoffs of external traits/actions) are implied by Lyapunov stability (which is implied by asymptotic stability).

Theorem 1′.

1. *If a state x^* is Lyapunov stable with recombination rate r, then it satisfies*
   - *(a) Weak internal stability: $\boldsymbol{w}^\intercal \cdot J^r_{x^*} \boldsymbol{w} \leq 0$ for every $\boldsymbol{w} \in T_{\mathrm{supp}(x^*)}$.*
   - *(b) External stability against traits: $u_{x^*} \geq u^r_{x^*}(a_d)$ for every $a_d \notin \mathrm{supp}_d(x^*)$.*
   - *(c) External stability against actions: $u_{x^*} \geq (1-r) u_{x^*}(\boldsymbol{a})$ for every $\boldsymbol{a} \notin \mathrm{supp}(x^*)$.*
2. *A stationary state x^* is asymptotically stable with recombination rate r if it satisfies the strict counterparts of conditions (a)–(c).*

Next we sketch the proof for Part (c) of Theorem 1′. The arguments for the remaining parts are similar to those of Theorem 1. See Appendix A.3 for the formal proofs.

Sketch of proof for Part (c) of Theorem 1′. Consider an external action $\hat{\mathbf{a}} \notin \mathrm{supp}(x^*)$. Consider the slightly perturbed state following an invasion of a few agents who play $\hat{\mathbf{a}}$: $(1-\varepsilon)x^* + \varepsilon\hat{\mathbf{a}}$. If $u_{x^*} < (1-r)u_{x^*}(\hat{\mathbf{a}})$, then $u_{x^*} < u^r_{x^*}(\hat{\mathbf{a}})$, which implies, due to Proposition 1′, that the share of agents playing action $\hat{\mathbf{a}}$ keeps growing, and hence that the population moves away from x^*. By contrast, if $u_{x^*} > (1-r)u_{x^*}(\hat{\mathbf{a}})$, then $u_{x^*} > u^r_{x^*}(\hat{\mathbf{a}})$ for a sufficiently small ε, which implies, due to Proposition 1′, that the share of agents playing action $\hat{\mathbf{a}}$ shrinks. □

Next we apply Theorem 1′ to characterise stable states in two interesting special cases: (I) pure states, and (II) interior states. Corollary 2′ shows that a pure state $\mathbf{a}$ is asymptotically stable essentially iff the payoff of $\mathbf{a}$ against itself is higher than (1) the payoff of any "neighbouring" action (which differs in a single trait) against $\mathbf{a}$, and (2) $1-r$ times the payoff of any action (which might differ in multiple traits).

Corollary 2′.

1. *If a pure state $\mathbf{a} \in A$ is Lyapunov stable with recombination rate r, then:*

   (a) $u_a(\mathbf{a}) \geq u_a(a'_d, \boldsymbol{a}_{-d})$ *for any dimension $d \in D$ and any trait $a'_d \in A_d$, and*

   (b) $u_a(\mathbf{a}) \geq (1-r) \cdot u_a(\mathbf{a}')$ *for any action $\mathbf{a}' \in A$.*

2. *A pure state $\mathbf{a} \in A$ is asymptotically stable with recombination rate r if it satisfies the strict counterparts of conditions (a)–(b).*

Example 1 (revisited). Corollary 2′ implies that any strict Nash equilibrium is asymptotically stable for every recombination rate r. Thus 1_m1_p is asymptotically stable for all values of r. The action 0_m0_p yields a higher payoff against itself than both neighbouring one-trait deviations, 0_m1_p and 1_m0_p. It is asymptotically stable if $5 = u_{0_m0_p}(0_m0_p) > (1-r)u_{0_m0_p}(1_m1_p) = 6(1-r) \iff r > \frac{1}{6}$, and it is unstable if $r < \frac{1}{6}$.

Corollary 3′ shows that an interior (full support) state x is asymptotically stable essentially if and only if the r-Jacobian matrix is negative definite.

Corollary 3′.

1. *If an interior state x^* is Lyapunov stable with recombination rate r, then* $\boldsymbol{w}^T J^r_{x^*} \boldsymbol{w} \leq 0$ *for every* $\boldsymbol{w} \in T_A$.

2. *An interior stationary state $x^* \in \Delta(A)$ is asymptotically stable with recombination rate r if $\boldsymbol{w}^T J^r_{x^*} \boldsymbol{w} < 0$ for every $\boldsymbol{w} \in T_A$.*

Corollaries 2′–3′ follow immediately from Theorem 1′.

6 Application to Organisational Decision-Making

This section applies the theoretical model to organisational decision-making. In the organisational interpretation, a player is an organisation, an action is a bundle of practices or decision components, and a trait is the component chosen in one domain.

Organisations frequently take complex decisions, such as setting strategic goals, that span multiple dimensions. Each dimension is typically overseen by a different division (e.g., marketing, product design, HR, logistics), with employees from backgrounds in marketing, engineering, and business. Moreover, it is rare for top management to have expertise in all areas, from technical production to business strategy (Gibbons and Henderson, 2013). The combinator dynamic is plausible when complex decisions are decomposed across subunits, so that adaptation is decentralised and component-wise rather than fully coordinated (see, e.g., Hortaçsu et al., 2024).

Next we revisit each of our three examples, and we present them with an organisational decision-making interpretation. Example 1 can be interpreted as a stylised model of complementary organisational change, where the two dimensions can refer to marketing (m) and product design (p). The payoffs capture the idea that adopting a single component of a new system may have little value, or may even reduce performance, while adopting the full bundle is valuable. This is a standard theme in the literature on complementarities in organisations (Brynjolfsson and Milgrom, 2013). Empirical work on information technology (Bresnahan et al., 2002), human-resource management (Ichniowski et al., 1997; MacDuffie, 1995), and enterprise-system implementation (Davenport, 1998; Scott and Vessey, 2002) provides concrete settings in which the value of one practice depends on the presence of complementary practices.

We interpret Example 2 as a mechanism for persistent performance differences. Each practice may appear locally acceptable when evaluated within its domain, even though the resulting bundle is low-performing because the components do not fit together (i.e., the marketing trait does not fit well with the product design trait). This interpretation is closely related to activity-system views of strategy (Porter, 1996) and to the evolution of fit among organisational choices (Siggelkow, 2002). It also speaks to the broad empirical fact that productivity and management performance differ persistently across firms and establishments (Syverson, 2011; Bloom and Van Reenen, 2007). Evidence from a field experiment in Indian textile plants shows that improved management practices can have large productivity effects and can diffuse slowly (Bloom et al., 2013). The example offers a dynamic explanation for why locally evaluated components need not eliminate performance dispersion.

Example 3 can be interpreted as a stylised model of specialised versus versatile organisational traits. A specialised practice may perform very well in one coherent configuration but poorly when recombined with mismatched components. A versatile or modular practice may perform reasonably well across a wider range of configurations. This interpretation is consistent with work on architectural innovation (Henderson and Clark, 1990) and modularity (Baldwin and Clark, 2000). In the model, component-wise recombination can make versatile traits attractive because they remain relatively robust when paired with different components. Whole-action imitation, by contrast, can sustain specialised high-performing bundles when coherent configurations can be recognised and copied.

The recombinator dynamics (Section 5) have a natural organisational interpretation as heterogeneity in coordination capacity. The whole-action imitation part of the dynamic corresponds to organisations that are able to recognise and adopt coherent bundles of practices. The trait-wise part corresponds to organisations in which learning is more decentralised and proceeds component by component. This interpretation is in the spirit of span-of-control and authority models (Lucas, 1978; Rosen, 1982) and is consistent with empirical evidence that managers affect organisational policies and performance (Bertrand and Schoar, 2003; Bandiera et al., 2020; Bennedsen et al., 2020).

7 Conclusion

This paper explores how multi-dimensional decision-making impacts strategic choices. It introduces the concept of combinator dynamics, where agents evaluate traits independently across different dimensions and combine them to form actions. This approach contrasts with traditional replicator dynamics, where agents imitate entire actions. The paper demonstrates that combinator dynamics can lead to different outcomes in terms of monotonicity, stationarity, and stability of states. It highlights that while traits with higher payoffs grow at a higher rate, this does not necessarily hold for multi-dimensional actions.

Through examples, the paper illustrates key insights of the combinator dynamics. First, the combinator dynamic can lead to suboptimal outcomes due to the difficulty of coordinated changes. Second, the combinator dynamic can produce stable states in which all traits have the same payoff even though the resulting actions have different payoffs. The application to organisational decision-making illustrates one setting in which the model is useful. Organisations often decompose complex decisions across domains, and local adoption of apparently successful practices can generate bundles whose components do not fit well together.

Our extended model of heterogeneous environments (recombinator dynamics) allows some players to imitate entire actions while others combine imitated traits. We show that in such environments, the stability of actions depends on the balance between the payoffs of actions and the frequency of traits. This induces a trade-off between specialisation and versatility of traits.

There is significant scope for future research to expand our theoretical framework in several intriguing directions. These include developing a multi-population version of the recombinator dynamics, establishing microfoundations through revision protocols, and exploring parallel models based on other evolutionary dynamics, such as best reply, logit, and Brown–von Neumann–Nash. Additionally, our current model assumes that the division of the action set into different dimensions is exogenous. It would also be of great interest to investigate a more comprehensive model where the dimensions are endogenous, shaped by the game's payoffs and the compatibility between different traits.

A Appendix

A.1 Proof of Proposition 1′

We rewrite the recombinator dynamics of Equation (4) in terms of the r-utility u^r_x defined in Equation (5):

$$\frac{\dot{x}(\mathbf{a})}{x(\mathbf{a})} = u^r_x(\mathbf{a}) - 1. \tag{11}$$

Part (1) of Proposition 1′ now follows immediately from Equation (11).

We turn to prove part (2). Fix any trait a_d in the support of x (i.e., $x(a_d) > 0$). Let us slightly rewrite Equation (4) from the perspective of a particular trait a_d:

$$\dot{x}(a_d, \mathbf{a}_{-d}) = (1-r)\frac{x(a_d, \mathbf{a}_{-d})u_x(a_d, \mathbf{a}_{-d})}{u_x} + r\frac{x(a_d)u_x(a_d)}{u_x^{|D|}}\prod_{d' \neq d} x(a_{d'})u_x(a_{d'}) - x(a_d, \mathbf{a}_{-d}). \tag{12}$$

Dividing Equation (12) by $x(a_d)$ and summing over all $\mathbf{a}'_{-d} \in A_{-d}$ yields the following *trait-centric recombinator dynamics*:

$$\begin{aligned} \frac{\dot{x}(a_d)}{x(a_d)} &= \sum_{\mathbf{a}'_{-d} \in A_{-d}} \frac{\dot{x}(a_d, \mathbf{a}'_{-d})}{x(a_d)} = \frac{(1-r)}{u_x}u_x(a_d) + \frac{r}{u_x^{|D|}}u_x(a_d)\sum_{\mathbf{a}'_{-d} \in A_{-d}}\prod_{d' \neq d} x(a'_{d'})u_x(a'_{d'}) - 1 \\ &= u_x(a_d)\left(\frac{(1-r)}{u_x} + \frac{r}{u_x^{|D|}}\sum_{\mathbf{a}'_{-d} \in A_{-d}}\prod_{d' \neq d} x(a'_{d'})u_x(a'_{d'})\right) - 1. \end{aligned} \tag{13}$$

Part (2) is now implied by the fact that replacing a_d with a'_d in the right-hand side of Equation (13) leaves the expression between the round brackets unchanged. □

A.2 Proof of Proposition 2′

Observe that (1) and (2) jointly imply that $u^r_x(\mathbf{a}) = 1$ for every $\mathbf{a} \in \text{supp}(x)$, which implies, by Proposition 1′, that x is stationary. For the other direction, suppose that x is stationary. By

Equation (11) from the proof of Part (1) of Proposition 1′ the stationarity of x implies that $u_x^r(\mathbf{a}) = 1$ for every $a \in \mathrm{supp}(x)$. By Part (2) of Proposition 1′ the stationarity of x implies that $u_x(a_d) = u_x$ for every d and $a_d \in \mathrm{supp}_d(x)$. Substituting this equality in the definition of $u_x^r(\mathbf{a})$ (Equation (5)) implies that $(1-r)\frac{u_x(\mathbf{a})}{u_x} + r\frac{\prod_{a_d \in a} x(a_d)}{x(\mathbf{a})} = 1$ for any $\mathbf{a} \in \mathrm{supp}(x)$.

A.3 Proofs of Theorem 1′

A.3.1 Proof Outline

The proof of Theorem 1′ proceeds as follows. We consider three different cases. The first case involves the dynamics of actions in the support of an equilibrium. The second case involves actions that have traits outside the support in multiple dimensions. The third and last case involves actions that have traits outside the support in a single dimension.

For each case we begin by considering the linear approximation of the dynamics near an equilibrium. In the first case we obtain a condition akin to that of the replicator dynamics, hence internal stability allows us to prove local stability. In the last case we show that the approximate dynamics is given by a diagonal matrix, and external stability for actions implies negative values along the diagonal.

The most challenging case is the second one. We first show that the linearisation of the dynamics there can be presented as linear dynamics over components whose partners have positive proportion in equilibrium plus some "external force" emerging from the interaction with actions with unsupported partners. It turns out that if the linear component over supported partners is negative definite then the dynamics as a whole is stable. To prove that it is indeed negative definite we introduce the notion of 'partner dynamics' and prove that these dynamics lead to a unique stable equilibrium which is in fact the stable partner distribution introduced in Equation (7). We then use this fact to show the negative definiteness of the original linearisation.

A.3.2 The Linear Approximation

We consider in this section a linear approximation of the recombinator dynamics around a stationary state x^*. The Lyapunov–Poincaré Theorem (Bellman, 1964, page 41) implies that the stability properties of the recombinator dynamics at a stationary state x^* will be implied by the stability properties of its linear approximation around the state x^*.

1. *Internal Stability*. Note that if one restricts attention to $\text{supp}(x^*)$ then the recombinator equation is a replicator-like equation with respect to the r-utility field u^r (in the sense that the form $\frac{\dot{x}(\mathbf{a})}{x(\mathbf{a})} = u^r_x(\mathbf{a}) - 1$ relates $\frac{\dot{x}(\mathbf{a})}{x(\mathbf{a})}$ on one side of the equation to a function of $\mathbf{a}$ on the other side). Hence expanding around x^* we obtain, since $\sum_{\mathbf{a}\in supp(x)} x(\mathbf{a})u^r_x(\mathbf{a}) = 1$, that (neglecting terms that are $O(||x - x^*||^2)$ for any $\mathbf{a} \in \text{supp}(x^*)$):

$$(14) \qquad \dot{x}(\mathbf{a}) = x^*(\mathbf{a}) \sum_{\mathbf{a}'\in\text{supp}(x)} \frac{\partial u^r_{x^*}(\mathbf{a})}{\partial x(\mathbf{a}')}(x(\mathbf{a}') - x^*(\mathbf{a}')) = x^*(\mathbf{a}) \sum_{\mathbf{a}'\in\text{supp}(x)} (J^r_{x^*})_{\mathbf{a},\mathbf{a}'}(x(\mathbf{a}') - x^*(\mathbf{a}')).$$

From Equation (14) it is immediately clear that it is sufficient for $J^r_{x^*}$ to be negative definite (with respect to the tangent space) for local asymptotic stability of the dynamics of $x(\mathbf{a})$ for $\mathbf{a} \in \text{supp}(x^*)$ to obtain. In the other direction, suppose that x^* is Lyapunov stable. Then $J^r_{x^*}$ must be negative-semi-definite: any positive value would initiate a trajectory that would eventually carry the state of the population away from x^*, contradicting Lyapunov stability.

2. *External stability against actions*. Next consider a type $\mathbf{a} \notin \text{supp}(x^*)$. Note that as the combinator part of the recombinator equation is a product of averages over traits a_d, the linear factors are obtained by actions $\mathbf{a}'$ such that $x^*(a'_d) = 0$ for a single value of d, namely: only for actions in which only one of their traits is outside the support of x^*.

If $\mathbf{a} \in A$ has multiple traits outside $\text{supp}(x^*)$, we obtain the approximate linear dynamic (neglecting terms that are $O(||x - x^*||^2)$)

$$(15) \qquad \dot{x}(\mathbf{a}) = x(\mathbf{a})\left((1-r)\frac{u_{x^*}(\mathbf{a})}{u_{x^*}} - 1\right).$$

It is clear from Equation (15) that it suffices for $(1-r)\frac{u_{x^*}(\mathbf{a})}{u_{x^*}} - 1 < 0$ for the weight of $\mathbf{a}$ to be reduced asymptotically (at an exponential rate), hence the same condition suffices for the local asymptotic stability of x restricted to neighbouring actions to hold. In the other direction, if $(1-r)\frac{u_{x^*}(\mathbf{a})}{u_{x^*}} - 1 > 0$, then Lyapunov stability cannot obtain, as the weight of $\mathbf{a} \notin \text{supp}(x^*)$ increases, carrying the population away from x^*.

3. *External Stability Against Traits*. We first consider the proof of Theorem 1′ for this part. Assume the dynamics is Lyapunov stable and that external-stability-against-traits does not hold, namely that $u^r_{x^*}(\hat{a}_d) - u_{x^*} > 0$ for some $\hat{a}_d \notin \mathrm{supp}_d(x^*)$. Define $\hat{\eta}^{\hat{a}_d}_{x^*} \in \Delta(A)$:

$$\hat{\eta}^{\hat{a}_d}_{x^*}(a) = \begin{cases} \eta^{\hat{a}_d}_{x^*}(\mathbf{a}_{-d}) & a_d = \hat{a}_d \text{ and } \mathbf{a}_{-d} \in \mathrm{supp}_{-d}(x^*) \\ 0 & \text{otherwise} \end{cases}$$

Consider a perturbed state $x = x_\varepsilon = (1-\varepsilon)x^* + \varepsilon\hat{\eta}^{\hat{a}_d}_{x^*}$ for sufficiently small $\varepsilon > 0$. Then $\hat{a}_d \in \mathrm{supp}_d(x)$ as $x(\hat{a}_d) = \varepsilon$, and

$$u_x(\hat{a}_d) = \frac{1}{x(\hat{a}_d)} \sum_{\mathbf{a}_{-d}\in A_{-d}} u_x(\hat{a}_d, \mathbf{a}_{-d}) x(\hat{a}_d, \mathbf{a}_{-d}) = \frac{1}{\varepsilon} \sum_{\mathbf{a}_{-d}\in A_{-d}} u_x(\hat{a}_d, \mathbf{a}_{-d}) \varepsilon \eta^{\hat{a}_d}_{x^*}(\mathbf{a}_{-d})$$

$$= \sum_{\mathbf{a}_{-d}\in A_{-d}} u_x(\hat{a}_d, \mathbf{a}_{-d}) \eta^{\hat{a}_d}_{x^*}(\mathbf{a}_{-d}).$$

Considering the last term, $u_x(\hat{a}_d)$ is sufficiently close to $u^r_{x^*}(\hat{a}_d)$ for every sufficiently small $\varepsilon > 0$ and therefore $u_x(\hat{a}_d) > u_{x^*}$. As $\hat{a}_d \in \mathrm{supp}_d(x)$ it follows from the monotonicity property proved in Proposition 1′ that, for every sufficiently small ε, if the trajectory is initialised at the perturbed state x_ε, namely $x^0 = x_\varepsilon$, then the weight $x^t(\hat{a}_d)$ increases monotonically, hence moving the population away from x^*, contradicting the assumption of Lyapunov stability.

The proof of local stability (part 2 of Theorem 1′) for this case is more involved and will be carried out in several stages. As before, we begin with the linear approximation for traits a_d in dimension d. Suppose that $\mathbf{a}_{-d} \in \mathrm{supp}_{-d}(x^*)$, while $(a_d, \mathbf{a}_{-d}) \notin \mathrm{supp}(x^*)$. The linear approximation for the equation of motion for $\mathbf{a}_{-d} \in \mathrm{supp}_{-d}(x^*)$ is (neglecting terms that are $O(||x - x^*||^2)$)

$$\text{(16)} \qquad \begin{aligned} \dot{x}(a_d, \mathbf{a}_{-d}) = (1-r)\frac{u_{x^*}(a_d, \mathbf{a}_{-d})}{u_{x^*}} x(a_d, \mathbf{a}_{-d}) \\ + \frac{r}{u_{x^*}} \prod_{d' \neq d} x^*(a_{d'}) \sum_{\mathbf{a}'_{-d}\in A_{-d}} u_{x^*}(a_d, \mathbf{a}'_{-d}) x(a_d, \mathbf{a}'_{-d}) - x(a_d, \mathbf{a}_{-d}). \end{aligned}$$

Note that if $\mathbf{a}'_{-d} \notin \text{supp}_{-d}(x^*)$ then

$$x^t(a_d, \mathbf{a}'_{-d}) = x^0(a_d, \mathbf{a}'_{-d}) e^{\left((1-r)\frac{u_{x^*}(a_d,\mathbf{a}'_{-d})}{u_{x^*}} - 1\right)t}, \tag{17}$$

where x^0 represents the initial state and x^t the state at time $t > 0$. Hence, setting

$$f_{a_d}(t) = \frac{r}{u_{x^*}} \prod_{d' \neq d} x^*(a_{d'}) \sum_{\mathbf{a}'_{-d} \notin \text{supp}_{-d}(x^*)} u_{x^*}(a_d, \mathbf{a}'_{-d}) x^0(a_d, \mathbf{a}'_{-d}) e^{\left((1-r)\frac{u_{x^*}(a_d,\mathbf{a}'_{-d})}{u_{x^*}} - 1\right)t}$$

We can rewrite Equation (16) for $\mathbf{a}_{-d} \in \text{supp}_{-d}(x^*)$ as

$$\begin{aligned} \dot{x}(a_d, \mathbf{a}_{-d}) = {} & (1-r)\frac{u_{x^*}(a_d, \mathbf{a}_{-d})}{u_{x^*}} x(a_d, \mathbf{a}_{-d}) \\ & + \frac{r}{u_{x^*}} \prod_{d' \neq d} x^*(a_{d'}) \sum_{\mathbf{a}'_{-d} \in \text{supp}_{-d}(x^*)} u_{x^*}(a_d, \mathbf{a}'_{-d}) x(a_d, \mathbf{a}'_{-d}) - x(a_d, \mathbf{a}_{-d}) + f_{a_d}(t), \end{aligned} \tag{18}$$

noting that $\int_0^\infty f_{a_d}(s) ds < \infty$.

Consider the matrix Λ^{a_d} indexed by elements $\mathbf{a}_{-d}, \mathbf{a}'_{-d} \in \text{supp}_{-d}(x^*)$ given by (which presents Equation (18) for a fixed a_d in matrix form):

$$\Lambda^{a_d}_{\mathbf{a}_{-d}, \mathbf{a}'_{-d}} = \begin{cases} \left((1-r) + r \prod_{d' \neq d} x^*(a_{d'})\right) \frac{u_{x^*}(a_d, \mathbf{a}_{-d})}{u_{x^*}} - 1, & \mathbf{a}_{-d} = \mathbf{a}'_{-d} \\ \frac{r u_{x^*}(a_d, \mathbf{a}'_{-d})}{u_{x^*}} \prod_{d' \neq d} x^*(a_{d'}), & \mathbf{a}_{-d} \neq \mathbf{a}'_{-d} \end{cases}. \tag{19}$$

Using Λ^{a_d}, and letting $f(t)$ be the vector whose entries equal $f_{a_d}(t)$, Equation (18) can be rewritten more compactly as

$$\dot{x}(a_d, \cdot) = \Lambda^{a_d} x(a_d, \cdot) + f(t), \tag{20}$$

where $\int_0^\infty ||f(s)|| ds < \infty$. The following lemma is then a direct corollary of Brauer (1964, Theorem 2):

Lemma 1. *If Λ^{a_d} is negative definite then all the solutions of Equation (20) tend to 0 as $t \to \infty$,* *and* 0 *is in fact Lyapunov stable.*

Proof. We prove Lemma 1 by applying (Brauer, 1964, Theorem 2). First we notice that Equation (20) has the form of (Brauer, 1964, Equation (1)). Indeed, in terms of (Brauer, 1964, Theorem 2) and (Brauer, 1964, Equations (5) and (9)), set $A = \Lambda^{a_d}$, $f(t, y) = f_{a_d}(t)$, and $p(t) \equiv 0$. Therefore, it is sufficient to prove that if $A = \Lambda^{a_d}$ is negative definite then $||e^{At}|| \leq Ke^{-\sigma t}$ for some $\sigma > 0$. This is an immediate corollary of the implication (3) $\Rightarrow$ (2) in Chicone (1999) (Chicone, 1999, Theorem 2.34). $\square$

Lemma 1 shows that to complete the proof we should show that Λ^{a_d} is negative definite. We will do that in the following section by introducing the notion of *partner distribution* and showing how it may be used together with the reduction above to prove stability.

A.3.3 The Partner Distribution and the Proof of the Main Results

We now begin to examine the negative definiteness of Λ^{a_d}. To prove this, it will be sufficient to prove that $x(a_d, \cdot) = 0$ is locally stable for the differential equation $\dot{x}(a_d, \cdot) = \Lambda^{a_d} x(a_d, \cdot)$, which by definition is the same as Equation (18) with the terms $f_{a_d}(t)$ neglected.

First, sum Equation (18), neglecting the terms $f_{a_d}(t)$, over $\mathbf{a}_{-d} \in \text{supp}_{-d}(x^*)$:

$$\dot{x}(a_d) = (1-r)\frac{\sum\limits_{\mathbf{a}_{-d} \in \text{supp}_{-d}(x^*)} x(a_d, \mathbf{a}_{-d}) u_{x^*}(a_d, \mathbf{a}_{-d})}{u_{x^*}} + r\frac{\sum\limits_{\mathbf{a}'_{-d} \in \text{supp}_{-d}(x^*)} u_{x^*}(a_d, \mathbf{a}'_{-d}) x(a_d, \mathbf{a}'_{-d})}{u_{x^*}} - x(a_d). \tag{21}$$

Consider the *partner distribution* $y(\mathbf{a}_{-d}) := \frac{x(a_d, \mathbf{a}_{-d})}{x(a_d)}$ for $\mathbf{a}_{-d} \in \text{supp}_{-d}(x^*)$, which is the ratio of the distribution of a_d to the weight of a_d at state x. Define the *normalised marginal trait-to-partner payoff* $u_{x^*}(a_d|y) = \sum\limits_{\mathbf{a}_{-d} \in \text{supp}_{-d}(x^*)} y(\mathbf{a}_{-d}) \frac{u_{x^*}(a_d, \mathbf{a}_{-d})}{u_{x^*}}$. From Equation (21) we obtain the differential equation:

$$\dot{x}(a_d) = (1-r)u_{x^*}(a_d|y)x(a_d) + ru_{x^*}(a_d|y)x(a_d) - x(a_d)$$
$$= (u_{x^*}(a_d|y) - 1)x(a_d) \tag{22}$$

Lemma 2. *Suppose that for every initial condition x^0 in an open neighbourhood of x^* the partner distribution $y^t \to \eta_{x^*}^{a_d}$ as $t \to \infty$ and furthermore $u_{x^*}(a_d|\eta_{x^*}^{a_d}) - 1 < 0$ as implied by external stability against traits. Then $x(a_d) = 0$ is a locally stable equilibrium of Equation (22).*

Proof. There is a $C > 0$ and time $t_0 > 0$ such that for every $t > t_0$ we have $u_{x^*}(a_d|y^t) - 1 < -C < 0$. Integrating Equation (22) we have, for some constant M and every $t > t' > t_0$

$$x^t(a_d) = Me^{\int_{t'}^{t}(u_{x^*}(a_d|y^s)-1)ds} < Me^{-C(t-t')}. \tag{23}$$

Taking $t \to \infty$ implies the lemma. □

We have thus shown that if the flow of the partner distribution y (as given by the solution of the differential equation $\dot{x}(a_d,\cdot) = \Lambda^{a_d}x(a_d,\cdot)$) converges to the stable partner distribution for every initial state near x^* and furthermore external stability against traits holds, then stability follows in this case as well.

In the last two sections of this appendix, we will show that the partner distribution follows a certain dynamic and that this dynamic has the stable partner distribution as a unique globally stable equilibrium, which will finish the proof of Theorem 1′.

A.3.4 Partner Dynamics

In this section we prove that the partner distribution, as defined by the solution of the differential equation $\dot{x}(a_d,\cdot) = \Lambda^{a_d}x(a_d,\cdot)$, satisfies a first-order differential equation. We call it the *partner dynamics*. In the next and final section we will prove that the stable partner distribution is the unique globally stable equilibrium of the partner dynamics.

Since $y(\mathbf{a}_{-d})x(a_d) = x(a_d,\mathbf{a}_{-d})$ one has $\dot{y}(\mathbf{a}_{-d})x(a_d) + y(\mathbf{a}_{-d})\dot{x}(a_d) = \dot{x}(a_d,\mathbf{a}_{-d})$. Dividing by $x(a_d)$ and rearranging we obtain

$$\dot{y}(\mathbf{a}_{-d}) = \frac{\dot{x}(a_d, \mathbf{a}_{-d})}{x(a_d)} - y(\mathbf{a}_{-d})\frac{\dot{x}(a_d)}{x(a_d)}. \tag{24}$$

Substituting this into Equation (18) (neglecting the $f_{a_d}(t)$ term), and considering the normalised payoff $\widehat{u}_{x^*}(a) = \frac{u_{x^*}(a)}{u_{x^*}}$ we obtain

$$\dot{y}(\mathbf{a}_{-d}) = (1-r)\widehat{u}_{x^*}(a_d, \mathbf{a}_{-d})y(\mathbf{a}_{-d}) + r\left(\prod_{d' \neq d} x^*(a_{d'})\right) u_{x^*}(a_d|y) - y(\mathbf{a}_{-d})u_{x^*}(a_d|y), \tag{25}$$

where, by definition $u_{x^*}(a_d|y) = \sum_{\mathbf{a}'_{-d} \in \text{supp}_{-d}(x^*)} \widehat{u}_{x^*}(a_d, \mathbf{a}'_{-d})y(\mathbf{a}'_{-d})$.

Going back to Lemma 2 in Appendix A.3.3, we are only left to prove that the partner dynamic is globally stable and that $y^t \to \eta^{a_d}_{x^*}$ as $t \to \infty$, where $\eta^{a_d}_{x^*}$ is the stable partner distribution. We prove this in the next section, Appendix A.3.5.

A.3.5 Global Stability of the Partner Dynamics

Recall that for a trait a_d and an associated partner distribution $y \in \Delta(\text{supp}_{-d}(x^*))$ we defined the partner dynamics as

$$\dot{y}(\mathbf{a}_{-d}) = (1-r)\widehat{u}_{x^*}(a_d, \mathbf{a}_{-d})y(\mathbf{a}_{-d}) + ru_{x^*}(a_d|y)\prod_{d' \neq d} x^*(a_{d'}) - y(\mathbf{a}_{-d})u_{x^*}(a_d|y), \tag{26}$$

where $\widehat{u}_{x^*}(\mathbf{a}) = \frac{u_{x^*}(\mathbf{a})}{u_{x^*}}$. Abusing notation (or alternatively, normalising the payoff units by the factor u_{x^*}) we will henceforth denote $u_{x^*}(a)$ instead of $\widehat{u}_{x^*}(a)$.

Note that η is a stationary point of these dynamics if and only if

$$(1-r)\frac{u_{x^*}(a_d, \mathbf{a}_{-d})\eta(\mathbf{a}_{-d})}{u_{x^*}(a_d|\eta)} + r\prod_{d' \neq d} x^*(a_{d'}) - \eta(\mathbf{a}_{-d}) = 0, \tag{27}$$

which means that stationary states of the partner dynamic of a trait correspond to the stable partner distribution as defined in Definition 4.

We now turn to prove the global stability proposition. First notice that the sign of the right-hand side of Equation (26) is the same as that of the left-hand side of Equation (27).

Fixing $\mathbf{a}_{-d}$, let η^0 be a distribution satisfying $\eta^0(\mathbf{a}_{-d}) = 0$, and for $t \in [0,1]$ set $\eta^t(\mathbf{a}_{-d}) = t$ and $\eta^t(\mathbf{a}'_{-d}) = (1-t)\eta^0(\mathbf{a}'_{-d})$ for $\mathbf{a}'_{-d} \neq \mathbf{a}_{-d}$. Plugging η^t instead of η into Equation (27) we obtain the equation

$$(1-r)\frac{tu_{x^*}(a_d, \mathbf{a}_{-d})}{tu_{x^*}(a_d, \mathbf{a}_{-d}) + (1-t)u_{x^*}(a_d|\eta^0)} + r\prod_{d' \neq d} x^*(a_{d'}) - t = 0. \tag{28}$$

Notice that the left-hand side of this last equation, which we denote $g_{\mathbf{a}_{-d}}(t|\eta^0)$ or $g(t)$ for short, is either a concave or a convex function of t. Furthermore, $g(0) = r\prod_{d' \neq d} x^*(a_{d'}) > 0$, and $g(1) = 1 - r + r\prod_{d' \neq d} x^*(a_{d'}) - 1 < 0$, which implies that g changes sign exactly once for every choice of distribution η^0.

Denoting the left-hand side of Equation (27) by $H_{\mathbf{a}_{-d}}(\eta)$, we can now deduce that the sets $V^+_{\mathbf{a}_{-d}} = \{\eta : H_{\mathbf{a}_{-d}}(\eta) > 0\}$ and $V^-_{\mathbf{a}_{-d}} = \{\eta : H_{\mathbf{a}_{-d}}(\eta) < 0\}$ are two open, disjoint, and connected sets. If $\eta(\mathbf{a}_{-d}) = 0$ we have $H_{\mathbf{a}_{-d}}(\eta) > 0$, and if $\eta(\mathbf{a}_{-d}) = 1$ we have $H_{\mathbf{a}_{-d}}(\eta) < 0$. We deduce that $|\eta^t(\mathbf{a}_{-d}) - \eta^{a_d}_{x^*}(\mathbf{a}_{-d})|$ decreases monotonically to 0 regardless of the initial point η^0, hence the dynamics satisfy the property of Lyapunov stability.

Furthermore, let η^* be a limit point of the dynamic, namely, suppose that there is a sequence $t_n \to \infty$ such that $\eta^{t_n} \to \eta^*$. Then it must hold that $H_{\mathbf{a}_{-d}}(\eta^*) = 0$ for every $\mathbf{a}_{-d} \in A_{-d}$, hence by the uniqueness of $\eta^{a_d}_{x^*}$ we have $\eta^* = \eta^{a_d}_{x^*}$.

References

Arad, A. and S. P. Penczynski (2024). Multi-dimensional reasoning in competitive resource allocation games: Evidence from intra-team communication. *Games and Economic Behavior 144*, 355–377.

Arad, A. and A. Rubinstein (2012). Multi-dimensional iterative reasoning in action: The case of the colonel blotto game. *Journal of Economic Behavior & Organization 84*(2), 571–585.

Arad, A. and A. Rubinstein (2019). Multidimensional reasoning in games: framework, equilibrium, and applications. *American Economic Journal: Microeconomics 11*(3), 285–318.

Arieli, I. and M. Mueller-Frank (2019). Multidimensional social learning. *The Review of Economic Studies 86*(3), 913–940.

Baldwin, C. Y. and K. B. Clark (2000). *Design Rules, Volume 1: The Power of Modularity*. Cambridge, MA: MIT Press.

Bandiera, O., A. Prat, S. Hansen, and R. Sadun (2020). CEO behavior and firm performance. *Journal of Political Economy 128*(4), 1325–1369.

Barton, N. H., S. Novak, and T. Paixão (2014). Diverse forms of selection in evolution and computer science. *Proceedings of the National Academy of Sciences 111*(29), 10398–10399.

Bellman, R. (1964). *Perturbation Techniques in Mathematics, Physics, and Engineering*. Holt and Winston.

Bennedsen, M., F. Pérez-González, and D. Wolfenzon (2020). Do CEOs matter? Evidence from hospitalization events. *Journal of Finance 75*(4), 1877–1911.

Bertrand, M. and A. Schoar (2003). Managing with style: The effect of managers on firm policies. *Quarterly Journal of Economics 118*(4), 1169–1208.

Bisin, A. and T. Verdier (2023). Advances in the economic theory of cultural transmission. *Annual Review of Economics 15*.

Bloom, N., B. Eifert, A. Mahajan, D. McKenzie, and J. Roberts (2013). Does management matter? Evidence from India. *Quarterly Journal of Economics 128*(1), 1–51.

Bloom, N. and J. Van Reenen (2007). Measuring and explaining management practices across firms and countries. *Quarterly Journal of Economics 122*(4), 1351–1408.

Börgers, T. and R. Sarin (1997). Learning through reinforcement and replicator dynamics. *Journal of Economic Theory 77*(1), 1–14.

Brauer, F. (1964). Nonlinear differential equations with forcing terms. *Proceedings of the AMS 15*(5), 758–765.

Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *Quarterly Journal of Economics 117*(1), 339–376.

Brynjolfsson, E. and P. Milgrom (2013). Complementarity in organizations. *The Handbook of Organizational Economics*, 11–55.

Chastain, E., A. Livnat, C. Papadimitriou, and U. Vazirani (2014). Algorithms, games, and evolution. *Proceedings of the National Academy of Sciences 111*(29), 10620–10623.

Chicone, C. (1999). *Ordinary Differential Equations with Applications*. Springer.

Cressman, R., A. Gaunersdorfer, and J.-F. Wen (2000). Evolutionary and dynamic stability in symmetric evolutionary games with two independent decisions. *International Game Theory Review 2*(01), 67–81.

Cressman, R. and Y. Tao (2014). The replicator equation and other game dynamics. *Proceedings of the National Academy of Sciences 111*(Supplement 3), 10810–10817.

Davenport, T. H. (1998). Putting the enterprise into the enterprise system. *Harvard Business Review 76*(4), 121–131.

Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press, Oxford, UK.

Edhan, O., Z. Hellman, and I. Nehama (2021). Making the most of potential: potential games and genotypic convergence. *Royal Society Open Science*.

Edhan, O., Z. Hellman, and D. Sherill-Rofe (2017). Sex with no regrets: How sexual reproduction uses a no regret learning algorithm for evolutionary advantage. *Journal of Theoretical Biology 426*, 67–81.

Eshel, I. and M. W. Feldman (1984). Initial increase of new mutants and some continuity properties of ess in two-locus systems. *The American Naturalist 124*(5), 631–640.

Frenkel, S., Y. Heller, and R. Teper (2018). The endowment effect as blessing. *International Economic Review 59*(3), 1159–1186.

Gibbons, R. and R. Henderson (2013). What do managers do? exploring persistent performance differences among seemingly similar enterprises. In R. S. Gibbons and J. Roberts (Eds.), *Handbook of Organizational Economics*, pp. 680–731. Princeton, NJ: Princeton University Press.

Hashimoto, K. (2006). Unpredictability induced by unfocused games in evolutionary game dynamics. *Journal of Theoretical Biology 241*(3), 669–675.

Henderson, R. and K. B. Clark (1990). Architectural innovation: the reconfiguration of existing product technologies and the failure of established firms. *Administrative Science Quarterly 35*, 9–30.

Herold, F. and N. Netzer (2023). Second-best probability weighting. *Games and Economic Behavior 138*, 112–125.

Hopkins, E. (2002). Two competing models of how people learn in games. *Econometrica 70*(6), 2141–2166.

Hortaçsu, A., O. R. Natan, H. Parsley, T. Schwieg, and K. R. Williams (2024). Organizational structure and pricing: Evidence from a large U.S. airline. *Quarterly Journal of Economics 139*(2), 1149–1199.

Ichniowski, C., K. Shaw, and G. Prennushi (1997). The effects of human resource management practices on productivity: A study of steel finishing lines. *American Economic Review 87*(3), 291–313.

Karlin, S. (1975). General two-locus selection models: some objectives, results and interpretations. *Theoretical population biology 7*(3), 364–398.

Lucas, R. E. J. (1978). On the size distribution of business firms. *Bell Journal of Economics 9*(2), 508–523.

MacDuffie, J. P. (1995). Human resource bundles and manufacturing performance: Organizational logic and flexible production systems in the world auto industry. *Industrial and Labor Relations Review 48*(2), 197–221.

Matessi, C. and C. Di Pasquale (1996). Long-term evolution of multilocus traits. *Journal of Mathematical Biology 34*(5), 613–653.

Maynard Smith, J. and G. R. Price (1973). The logic of animal conflict. *Nature 246*(5427), 15–18.

Meir, R. and D. C. Parkes (2015). On sex, evolution, and the multiplicative weights update algorithm. In *Proc. of the International Conference on Autonomous Agents & Multiagent Systems*.

Mertikopoulos, P. and W. H. Sandholm (2018). Riemannian game dynamics. *Journal of Economic Theory 177*, 315–364.

Netzer, N., A. Robson, J. Steiner, and P. Kocourek (2021). *Endogenous Risk Attitudes*. Centre for Economic Policy Research.

Newton, J. (2018). Evolutionary game theory: A renaissance. *Games 9*(2), 31.

Palaiopanos, G., I. Panageas, and G. Piliouras (2017). Multiplicative weights update with constant step-size in congestion games: convergence, limit cycles and chaos. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 5874–5884.

Porter, M. E. (1996). What is strategy? *Harvard Business Review 74*(6), 61–78.

Read, D., G. Loewenstein, M. Rabin, G. Keren, and D. Laibson (2000). Choice bracketing. *Elicitation of preferences*, 171–202.

Rosen, S. (1982). Authority, control, and the distribution of earnings. *Bell Journal of Economics 13*(2), 311–323.

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

Sawa, R. and D. Zusai (2014). Evolutionary imitative dynamics with population-varying aspiration levels. *Journal of Economic Theory 154*, 562–577.

Sawa, R. and D. Zusai (2019). Evolutionary dynamics in multitasking environments. *Journal of Economic Behavior & Organization 166*, 288–308.

Scott, J. E. and I. Vessey (2002). Managing risks in enterprise systems implementations. *Communications of the ACM 45*(4), 74–81.

Siggelkow, N. (2002). Evolution toward fit. *Administrative Science Quarterly 47*(1), 125–159.

Steiner, J. and C. Stewart (2016). Perceiving prospects properly. *American Economic Review 106*(7), 1601–31.

Syverson, C. (2011). What determines productivity? *Journal of Economic Literature 49*(2), 326–365.

Taylor, P. D. and L. B. Jonker (1978). Evolutionary stable strategies and game dynamics. *Mathematical Biosciences 40*(1-2), 145–156.

Thaler, R. H. (1999). Mental accounting matters. *Journal of Behavioral Decision Making 12*(3), 183–206.

Van Damme, E. (1991). *Stability and perfection of Nash equilibria*, Volume 339. Springer.

Waldman, M. (1994). Systematic errors and the theory of natural selection. *The American Economic Review*, 482–497.

Weibull, J. W. (1997). *Evolutionary Game Theory*. Cambridge, MA: MIT Press.

Williams, G. C. (1966). *Adaptation and natural selection*. Princeton, NJ: Princeton University Press.