





Social decision-making driven by artistic explore–exploit tension

Kayhan Özcimder^a, Biswadip Dey^a, Alessio Franci ^b, Rebecca Lazier^c,
Daniel Trueman^d and Naomi Ehrich Leonard ^a

^aDepartment of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ, USA;

^bDepartment of Mathematics, National Autonomous University of Mexico, Mexico City, Mexico; ^cProgram in Dance, Lewis Center for the Arts, Princeton University, Princeton, NJ, USA; ^dDepartment of Music, Princeton University, Princeton, NJ, USA

ABSTRACT


We studied social decision-making in the rule-based improvisational dance *There Might Be Others*, where dancers make in-the-moment compositional choices. Rehearsals provided a natural test-bed with communication restricted to non-verbal cues. We observed a key artistic explore–exploit tension in which the dancers switched between exploitation of existing artistic opportunities and riskier exploration of new ones. We investigated how the rules influenced the dynamics using rehearsals together with a model generalized from evolutionary dynamics. We tuned the rules to heighten the tension and modelled nonlinear fitness and feedback dynamics for mutation rate to capture the observed temporal phasing of the dancers' exploration-versus-exploitation. Using bifurcation analysis, we identified key controls of the tension and showed how they could shape the decision-making dynamics of the model much like turning a 'dial' in the instructions to the dancers could shape the dance. The investigation became an integral part of the development of the dance.


KEYWORDS

Dance improvisation;
explore–exploit decision-
making; communication
through movement;
replicator-mutator and
co-evolutionary dynamics;
collective behaviour;
pitchfork bifurcation

1. Introduction

Social decision-making enables group efforts that are neither fully scripted nor centrally controlled (Dyer et al. 2009), so predicting group behaviour requires understanding what drives individual choices (Sanfey 2007). Social decision-making is studied in various research communities, including social choice theory (Arrow 1951; Sen 1970; Saari 2001), which combines social ethics with voting theory, social neuroeconomics (Fehr and Camerer 2007; Sanfey 2007; Lee 2008), which joins game theory with psychology and neuroscience, and collective animal behaviour (Seeley and Buhrman 1999; Conradt and Roper 2003; Couzin et al. 2005; Meunier et al. 2006; Leonard et al. 2012). In network science, researchers examine the role of network structure in decentralized decision-making groups (Lazer and Friedman 2007; Olfati-Saber, Fax, and Murray

CONTACT Naomi Ehrich Leonard  naomi@princeton.edu

 Supplemental data for this article can be accessed at <https://doi.org/10.1080/03080188.2018.1544806>

© 2018 Institute of Materials, Minerals and Mining Published by Taylor & Francis on behalf of the Institute

2007; Mason and Watts 2009; Greening, Pinter-Wollman, and Fefferman 2015; Landgren, Srivastava, and Leonard 2016).

A fundamental consideration in settings where options yield rewards is how decision-makers make choices that balance exploitation of options with well-known rewards with the riskier, but possibly advantageous, exploration of options with poorly known rewards. In reward-based decision-making, the decision-maker chooses from a set of options and receives a reward associated with the chosen option. Rewards may be uncertain and variable, and so the decision-maker seeking to maximize reward faces a dilemma between choosing an option that is known to yield reasonably high reward (exploitation) versus choosing an option for which they have little information but that could yield even higher reward (exploration). Exploitation and exploration are in tension because a decision-maker who only exploits will not get the new information that comes from exploring, and a decision-maker who only explores will not leverage the new information received.

When the decision-maker is allowed a sequence of reward-bearing choices over time, the explore–exploit tension changes with each choice. For example, consider choosing a restaurant for dinner in a city where you are visiting for an extended time. Suppose you have an outstanding meal at restaurant A on the first night. Exploiting the newly gained information, you return the second night but have merely a decent meal. Then, on the third night, you wonder if you should exploit your two data points and go back to restaurant A for what you can anticipate will be something in between decent and outstanding. Or do you explore a restaurant B hoping to find a consistently outstanding meal?

A large literature addresses the explore–exploit tension for single decision-makers and provides algorithms for choices that reliably optimize the accumulation of reward over time (Lai and Robbins 1985; Auer, Cesa-Bianchi, and Fischer 2002; Kaufmann, Cappé, and Garivier 2012). The tension is examined in a wide range of contexts, including control of attention in the brain (Cohen, McClure, and Yu 2007), allocation of treatments in clinical trials (Villar, Bowden, and Watson 2015), and animal foraging in a patchy environment (Krebs, Kacelnik, and Taylor 1978). The tension has more recently been studied for groups of interacting decision-makers, using simulation (Lazer and Friedman 2007), experiment (Mason and Watts 2009), and model-based analysis (Kalathil, Nayyar, and Jain 2014; Kolla, Jagannathan, and Gopalan 2016; Landgren, Srivastava, and Leonard 2016). In these group decision-making settings, individuals make their own choices among options, but they are influenced by social interactions that involve observations and communications of the choices or evaluations that others in the group make.

The studies of decision-making groups raise important open questions about how factors in the social interactions influence the explore–exploit tension and the group's decision-making dynamics. How important is the clarity of the communication of choices and rewards? Does it matter if interactions are with experts or with novices? What is the role of the network structure of the interactions, i.e. who can observe or communicate with whom? Mason and Watts (2009) ran an experiment in which each participant was asked to make a sequence of choices of location on a computer generated map in order to find the deepest oil well. After every choice, each participant could see their reward (how deep the oil at the location selected) as well as the choices and rewards of some of the other participants. The data showed that networks (as compared to solitary decision-making) suppressed individual exploration, i.e. individuals copied one another

more often than trying out a new location. However, when someone in the group made a highly rewarding choice, the whole group benefited.

The authors cautioned that their results were likely dependent on the structure of the environment, where there was only one location with a deep oil well and thus little information to be gained by exploring the rest of the map. The average reward at each location also remained constant, which has been the case in most settings examined in the literature. When average rewards do not change, decision-makers can ultimately stop exploring once they have enough information to find the best option. These observations illustrate the importance of understanding the influence of the environment on the explore–exploit tension and the group’s decision-making dynamics. For instance, if the rewards change, how frequently will the decision-maker explore? What happens if environmental conditions create urgency in decision-making? How does decision-making change if sticking with the current option is easy and choosing a new one is costly? What if the reward associated with an option decreases if a decision-maker sticks with it too long? Or what if the reward is diminished when too many individuals select the same option at the same time?

In the present paper, we report on a novel, generative investigation of *There Might Be Others* (TMBO), an open choreographic work where performers collectively create the piece in real-time negotiating a catalogue of defined movement ‘modules’ with a set of performer instructions and governing rules. TMBO builds on the tradition of open scores and improvisational works wherein the performers compose the work in performance within a set of rules and contingencies. The artistic quality of TMBO unfolds as the dancers experiment with relationships, timing, space, and groupings and collectively work to ensure unpredictability. Through their decisions balancing how, when, and where to perform the defined vocabulary, the dancers create beautiful moments of juxtaposition, complex groupings, and dynamic shifts in tempo.

Our investigation arose out of an art and science collaboration aimed at finding principled ways to influence the creative process in social, rule-based art-making. As it turned out, the investigation led to this and much more, including a fresh perspective on the development of the piece and the means to examine open scientific questions. TMBO provided a rich opportunity for studying social decision-making in a creative endeavour by highly trained artists, with rehearsals serving as a natural test-bed and dynamics well-suited to mathematical modelling and analysis. In TMBO, the dancers’ decisions carry both artistic rewards and risks. Recognizing this, we identified an explore–exploit tension driving the artistic choices of the dancers, and a connection between the tension and the rules of the dance. We defined choosing to ‘exploit’ as joining a module currently being danced, and choosing to ‘explore’ as introducing a module that has not yet been danced. These are in tension, and yet the resulting dynamics are quite different from the standard explore–exploit dynamics in which average rewards are stationary. In TMBO, rewards can change depending on the sequence of choices, the changing environment, and the dancers’ artistic sensibilities. For example, a module that initially was highly rewarding might lose its appeal if it is danced for too long; one or more dancers might then seek to explore something new. The result is a richly varied dynamic involving periodic switching between exploration and exploitation.

To investigate the mechanisms at play and the opportunities for design, we made a systematic examination of the rules, environmental context, explore–exploit tension and the

overall effect on the decision-making dynamics and the dance. Our approach integrated ongoing rehearsals with the TMBO dancers and analysis of a low-dimensional mathematical model of the explore–exploit tension, which we generalized from evolutionary dynamics (Bürger 1998). In rehearsal, we modified rules and recorded the resulting behaviour of the dancers. With the model we studied the sensitivity of the explore–exploit dynamics with respect to parameters in the modelled rules and context. Our observations in rehearsal led to refinements in the model. And the model provided an off-line tool for examining mechanisms and for investigating new opportunities before trying them in rehearsal and designing them into practice. We learned much about the driving forces behind the underlying switching dynamics of the dance and how these dynamics could be shaped with subtle modification to instructions. Along the way, the language of our explore–exploit study became our baseline language for making TMBO.

In this paper, we describe our investigation of the artistic explore–exploit tension in TMBO, its driving influence on the group’s decision-making dynamics, and the unfolding of the dance. We reveal how the investigation itself became an integral part of the development of the dance. And we show how the results of the investigation, including the dance and the model, provide new ways of understanding and shaping these and related social decision-making dynamics.

2. Structure and rules of the dance

TMBO was conceived by choreographer Rebecca Lazier who was inspired by the musical piece *In C* composed by Terry Riley in 1964. Mid-development, Rebecca invited composer Dan Trueman and engineering professor Naomi Leonard to join the project; they were examining emergence in structured improvisation and were inspired by both *In C* and the art-science project *Flock Logic* (Leonard et al. 2014). In *In C*, musicians make compositional choices as they play an ordered sequence of 53 musical melodic patterns (Riley 1964; Reed 2011). All of these patterns fit on a single page, and the instructions for the piece are similarly concise. *In C* is both participatory and emergent, relying on collective musical decision-making from a group of expert performers to create a performance that is always recognizable but also unique; it is widely considered one of the most influential pieces of twentieth-century concert music, and the beginning of musical Minimalism (Carl 2009).

Similarly, TMBO is defined by a catalogue of 44 movement modules (in analogy to the musical melodic patterns) and performing instructions that lay out the choreographer’s artistic objectives and rules for how a group of dancers can manipulate the modules in performance (Lazier et al. 2016). The objectives and rules also apply to how a group of dancers can order the modules in performance, since, unlike in the Riley piece, no prescribed order is set in the TMBO score. Riley instructs the players to repeat each melodic pattern as many times as they wish before moving on to the next, to proceed in order, and to listen to the group to stay within three patterns of one another. In TMBO, performers use these general rules, except that dancers are responsible for the compositional tone of the work by selecting which module should come next based on the needs of the piece in the moment. Once a module is introduced in the dance, the order is then set forth and followed by the entire cast. However, no one can lead two modules in a row, as the leadership must be shared by the group. Using rules and

games to create dance has a long and rich history (Clemente 1990). There is also a growing body of research on cognition and the distributed choreographic process (Kirsh 2011). In Guimerà et al. (2005), the mechanisms of self-assembly of teams in creative endeavours are examined.

TMBO's 44 modules include a balance of composed movement, gestural ideas, and tasks. Some modules are boisterous, such as *Jump Bean*, where the instruction is to bounce and interact with fellow dancers, others meditative, such as *Oar*, where the body is gently shifting forward and backward while one hand swoops in a paddle-like action. There are complex phrases where performers will join in unison, jumping patterns that slice through the space, and single actions like *Human Pile* where performers pile on top of each other in stillness. The aesthetic influences range from ballet, contemporary, and modern dance to hip hop and various folk forms. The modules are the vocabulary of TMBO and while each is defined and recognizable, they can be adapted according to rules for each module in order to fulfil the compositional needs of the piece in the moment. In devising the modules, we learned that they need to have dynamic, rhythmic, expressive, spatial, and stylistic range to support an aesthetic of the work that is multifaceted and unpredictable. Each dancer must memorize the catalogue of modules and be able to not only select the one most suited for the moment but also be willing to continually develop the performance of the modules in new ways with each iteration. Figure 1 shows two photos from a performance of TMBO at New York Live Arts, New York City, in March 2016.

TMBO starts with dancers gradually joining the performance of a single module and proceeds with dancers introducing modules. As a new module is introduced the order is then set and each performer must proceed in the order modules appeared; in the general rules they cannot skip or switch modules. Any dancer can introduce a new module at any time, although no more than one module can be introduced at the same time nor can a dancer lead two modules in a row; it is imperative that leadership shift among the entire population. Once a module has been introduced, danced, and abandoned, it cannot be re-introduced. Modules can be juxtaposed by different subgroups of dancers performing different modules simultaneously. However, the dancers are restricted to a limited number of modules at a time; in the general rules, the limit is three modules at a time. Dancers can experiment with and innovate variations of any current module at any time, as long as they respect the rules and keep modules recognizable to the cast to ensure clear communication. The dance is over after a fixed time or when all the modules have been performed, whichever comes first. Because none of the dancers' decisions are made in advance and they are instructed to never repeat what they have done before, each rehearsal and performance brings fresh choices of sequencing, timing, counterpoint, relationships, and variation of modules.

By design, the dancers retain considerable creative freedom: they choose which and when, as well as where and how, modules are performed to meet artistic goals for juxtaposition, unpredictability, and dynamic pacing. Social interactions are a priority: the dancers should clearly communicate their choices through their movement and other non-verbal cues, keep up frequent observation of the movement choices of the other dancers, and make their own choices in response to what the other dancers are doing. For example, one artistic objective seeks texture and richness through juxtaposition of modules, and this can only be achieved through coordinated choices among the



Figure 1. Photos from a TMBO performance at New York Live Arts, New York City, in March 2016. In the top photo, several dancers on the left side perform the module *Flamingo* and several dancers on the right side perform the module *Folk*, while a pair in the foreground perform the module *Box and Drag*. In the bottom photo, three dancers on the sides perform the module *Arms* and two dancers in the centre perform the module *Jump Bean*. Photo credit: Ian Douglas.

dancers, such as when a fraction of the group chooses to perform *Arms* while the rest performs *Jump Bean* as in the bottom photo of Figure 1. Another artistic objective seeks recurring moments of surprise, each defined by the introduction of a new module. This too can only be achieved through coordinated choices among dancers, since only one dancer can introduce a module at a time, and the limit on the number of current modules cannot be exceeded.

Through direct interactions with and observations of dancers rehearsing and performing TMBO, we discovered that the social decision-making during TMBO is driven in large part by an explore–exploit tension in which rewards are artistic in nature. By choosing to

experiment with an existing module, the dancer exploits a known option. And by choosing to introduce a new module, the dancer takes the riskier step and explores a less predictable option. Both are creative choices and can add artistically to the performance. We found that a key to how the dancers trade off exploiting and exploring is the performance rule that limits the number of modules that should be danced at a time. Originally, a limit of three modules was loosely applied, meaning that the dancers would sometimes add a fourth module. However, if the limit is made strict and there are already three modules being danced, a dancer cannot add a new module until all the dancers coalesce into two modules. In this way, exploiting and exploring are strongly in tension for the group: *either* the dancers experiment with all the current modules *or* they complete a current module so that one dancer can introduce a new module. The tension exists in principle for any limit on number of current modules; however, we found that the lower the limit the greater the tension. So for our investigation, we set the strict limit to be two modules at a time, which means a dancer can add a new module only when all dancers have converged on a single module.

The musical component of TMBO operates similarly to the dance. The musicians have their own set of modules, and while they are conceptually paired with the dance modules, in performance the musicians set their own module order and may not be lined up with the dancers. A full exploration of the musical component of TMBO is beyond the scope of this paper, and our research focus here was on the dance; in the future, it would be of interest to pursue a similar exploration of the musical decision-making dynamics, and how the musical and dance dynamics might interact.

3. Studio rehearsals

We made a systematic investigation of the influence of changes in performance rules and context on the dancers' decision-making dynamics and the dance, during a rehearsal in the Patricia and Ward Hagan '48 Studio, Lewis Center for the Arts, Princeton University on 25 July 2015. Because the aim was to find principled ways to modify the structure to influence the choreographic development of the dance, the Institutional Research Board (IRB) at Princeton University did not require that we seek IRB approval for the study reported in this paper.

Our systematic study focused on three abbreviated run-throughs of TMBO, each lasting 9–10 minutes, rather than the full 60 minutes of a TMBO performance. The catalogue of modules available to the dancers was also reduced from 44 to 9 modules. The dancers were instructed to start with the module *Clapping*, which was one of the nine in the catalogue for each run-through. Nine professional dancers participated in the rehearsal: Rhonda Baker, Simon Courchel, Natalie Green, Raja Feather Kelly, Cori Kresge, Christopher Ralph, Tan Temel, Saúl Ulerio, and Shayla Vie-Jenkins. All nine dancers took part in Run-through 1, but only eight of the dancers took part in Run-throughs 2 and 3 since Saúl Ulerio had to leave early. For all three run-throughs we fixed the following performance rules, slightly modified from the general rules:

- 1 any dancer can introduce a new module at any time
- 2 a module cannot be chosen if it has come and gone
- 3 any dancer can switch to any current module at any time

- 4 any dancer can skip a module
- 5 no more than two modules can be danced at a time.

Rules (1) and (2) are unchanged from the general rules: only one module can be introduced at a time and the same dancer should not introduce two modules in a row. Rules (3) and (4) are new relative to the general rules where the set order is to be followed. Rule (5) is a modification of the original loose enforcement of a maximum of three modules at a time to a strict enforcement of a maximum of two modules at a time.

Before they performed the run-throughs, the dancers were not informed of the motivation for the modified rules, the model, nor the investigation of an explore–exploit tension. Each run-through was recorded with two high-definition video cameras angled to capture the whole stage.

3.1. Run-through 1

In Run-through 1, the catalogue of available modules included the following nine: *Clapping*, *Flamingo*, *Chain*, *Whip It*, *Crawl and Sing*, *Trigger*, *Do Op*, *Lay Down and Get Up*, *Drop and Roll*. The run-through began with dancers starting the module *Clapping*.



Figure 2. Run-through 1.

A snapshot at 37 s into Run-through 1 is shown in [Figure 2\(a\)](#). In the snapshot, three dancers, Simon, Natalie, and Tan, performed *Clapping* facing one another in the centre of the space; Rhonda and Raja can be seen on the edges of the rehearsal space not participating in *Clapping*, although Raja had briefly been clapping.

In the snapshot of [Figure 2\(b\)](#), at 42 s into Run-through 1, Cori can be seen walking and clapping. Natalie, who is lying down on her back, has clearly introduced the module *Lay Down and Get Up*. Since there was only one active module, *Clapping*, any dancer could have decided to introduce a new module among the remaining eight in the catalogue at any moment, but Natalie chose *Lay Down and Get Up*. At this point in time, since there were two active modules, *Clapping* and *Lay Down and Get Up*, no new module could have been introduced.

In the snapshot of [Figure 2\(c\)](#), at 69 s into Run-through 1, the dancers were distributed over the two active modules, *Clapping* and *Lay Down and Get Up*, although Natalie had switched back to *Clapping*. Cori, Chris, Rhonda, Saúl, and Tan were also performing *Clapping*. Note that Chris (second from the left) had begun to modify *Clapping*; here he can be seen clapping with his arms rather than his hands. Natalie was also modifying *Clapping* by moving her shoulders to the beat. Shayla, lying on her side, Simon, lying on his stomach, and Raja, lying on his back, were all performing *Lay Down and Get Up*.

At 146 s into Run-through 1, when all the dancers coalesced, at least momentarily into *Lay Down and Get Up*, Cori introduced the module *Whip It* and Shayla almost immediately joined her. The snapshot of [Figure 2\(d\)](#) shows the dancers at just this moment, with Cori and Shayla performing *Whip It* and the rest of the dancers evolving *Lay Down and Get Up*. Since *Clapping* had come and gone, the dancers were no longer allowed to perform it for the remainder of Run-through 1.

By 275 s into Run-through 1, eight of the dancers were performing *Whip It* and only one dancer (Rhonda) remained with *Lay Down and Get Up*. A snapshot at 298 s into Run-through 1 is shown in [Figure 2\(e\)](#). This lasted until around 335 s when Rhonda finally joined *Whip It*. The module *Do Op* was then immediately introduced. *Whip It* was over shortly thereafter and *Drop and Roll* was introduced. A snapshot at 412 s into Run-through 1 is shown in [Figure 2\(f\)](#), where six dancers can be seen performing *Do Op* and three dancers can be seen (mid-roll on the floor) performing *Drop and Roll*. The dancers remained with *Do Op* and *Drop and Roll* until Run-through 1 ended at 540 s (nine minutes). By the end, all nine dancers were performing *Do Op*. Although they were available to the dancers, the four modules *Flamingo*, *Chain*, *Crawl and Sing*, and *Trigger* were not selected during Run-through 1.

3.2. Run-through 2

In Run-through 2, we added to the instructions of Run-through 1 by asking the dancers to be more impulsive in their choices, i.e. to *increase* their tendency to switch modules. The intent was to examine switching tendency as a possible design ‘dial’ that could be cued to the whole group of dancers during a run-through, for example, by signaling with the beat of a drum. Switching tendency was identified as a candidate design dial since it was revealed to be associated to critical parameters in the analysis of the mathematical model as described below. For Run-through 2 the catalogue of available modules included

the following nine: *Clapping*, *Arms*, *Box*, *Bravo*, *Bumper Body*, *Chris*, *Folk*, *Kiss*, *Romper Room*.

Run-through 2 began with dancers starting the module *Clapping*. A snapshot at 38 s into Run-through 2 is shown in Figure 3(a). In the snapshot, three dancers, Simon, Shayla, and Raja, can be seen performing *Clapping*. Raja had already begun modifying *Clapping* by clapping his hands on his knees. The module *Bumper Body* was introduced and the snapshot at 109 s in Figure 3(b) shows Cori, Chris, Natalie, and Tan, all performing *Bumper Body* while Shayla, Simon, and Raja continued in *Clapping*.

When *Bumper Body* ended, *Folk* was introduced. When *Folk* ended *Romper Room* was introduced and when *Romper Room* ended *Chris* was introduced. All the while *Clapping* persisted. And dancers frequently switched back and forth between *Clapping* and whichever other module was current. For example, at 157 s in Figure 3(c), Natalie, Simon, Chris, and Shayla were all performing *Folk*. But at 228 s in Figure 3(d), Shayla and Cori were the only two dancers performing *Folk*, and at 251 s in Figure 3(e), Raja was alone doing *Folk*. Those dancers who had returned to *Clapping* were arranged in a line, and it was if the dancers were taking turns doing solo performances of *Folk*.

All of the dancers but Chris can be seen performing the module *Chris* at 368 s in Figure 3(f). Chris was keeping *Clapping* alive. In fact, *Clapping* had another resurgence of



Figure 3. Run-through 2.

popularity before it was abandoned and *Bravo* introduced. At the 9-minute mark, one dancer remained in *Chris* while the rest performed *Bravo*.

3.3. Run-through 3

In Run-through 3, the catalogue of nine modules was the same as in Run-through 2. However, the additional rule was changed: in Run-through 3, we asked the dancers to be more resistant to changing their choices, i.e. to *decrease* their tendency to switch modules. The intent was to continue to examine switching tendency as a possible design dial. However, in contrast to Run-through 2 where the dial was turned up, in Run-through 3 the dial was turned down.

Run-through 3 began with the dancers starting the module *Clapping*. At 70 s into the run-through, the module *Box and Drag* was introduced. *Box and Drag* persisted while *Clapping* ended and *Kiss* was introduced, *Kiss* ended and *Folk* was introduced, *Folk* ended and *Bravo* was introduced. *Box and Drag* was still active at the 9-minute mark, along with *Bravo*.

Throughout Run-through 3, there was a minimal switching between active modules; for example, there was nothing like the switching of soloists performing *Folk* as seen in Run-through 2. When Shayla introduced *Folk* at 286 s into Run-through 3, she performed it alone for 48 s before she was joined by Simon, and then some of the others. At 450 s, Chris rose up from *Box and Drag* and caught each dancer who was performing *Folk* and folded the dancer into *Box and Drag* until at 490 s, he managed to get everyone committed to *Box and Drag* and *Folk* was gone. At just this moment, Chris introduced *Bravo*. Later, after the 9-minute mark, when almost all the dancers were performing *Bravo*, Cori seemed to play with the idea of folding some of them into *Box and Drag*, something like an echo of what we had seen Chris do earlier.

After Run-through 3, we ran one more run-through during which we tried out a global cue to signal a change in instruction, i.e. to signal a dial change. One beat on the drum directed the dancers to make more ‘studied’ choices, whereas a scale on the xylophone directed the dancers to make more ‘random’ choices. The dancers were a bit worn out at this point, but they described enjoying the exercise. Raja described liking the connection to the external environment rather than only to internal, i.e. social, factors. Cori reported liking the global external signal because it served to refresh in her mind a priority.

3.4. Quantifying group dynamics

Motivated by our modelling approach, described below, we quantified the explore–exploit dynamics of the group of dancers in terms of the distribution of dancers over the modules. For Run-throughs 1, 2, and 3, we visualized the changing distribution by plotting in [Figure 4](#) the fraction of total number of dancers performing each of the current modules as a function of time. Videos extracted from the cameras were synchronized using off-the-shelf video editing tools (*SI Appendix*, Figure S1). The synchronized videos were then analysed by counting the number of dancers for each one-second time period and the trajectories drawn by connecting the sample points.

The top, middle, and bottom plots of [Figure 4](#), show what fraction of the group of dancers in Run-through 1, 2, and 3, respectively, performed which modules, distinguished




Figure 4. The fraction of total number of dancers performing each of the at most two modules, x_1 and x_2 as a function of time t during Run-through 1 (top plot), Run-through 2 (middle plot), and Run-through 3 (bottom plot). For Run-through 1, the x_1 and x_2 are fractions of a total of nine dancers performing each of the two active modules, whereas for Run-throughs 2 and 3, they are fractions of a total of eight dancers. The modules are distinguished by the colours indicated in the key on the right.

by colour, over time. Time, on the horizontal axis, is denoted by t and expressed in seconds from the start of the run-through at $t=0$ s until the 9-minute mark at $t=540$ s.

We denote the fraction of dancers at time t performing one of the two current modules by $x_1(t)$ and the fraction of dancers at time t performing the other of the two current modules by $x_2(t)$. The fractions $x_1(t)$ and $x_2(t)$ therefore take values between 0 and 1 for all t . Since only two modules were allowed at a time, $x_1(t) + x_2(t) = 1$ if all dancers were dancing at time t . For instance, at $t=300$ s into Run-through 1, reading off the higher (green) line in the top plot of Figure 4, the fraction of dancers performing *Whip It* was $x_1(300) = 8/9$ (everyone but Rhonda) and reading off the lower (red) line, the fraction of dancers performing *Lay Down and Get Up* was $x_2(300) = 1/9$ (Rhonda).

Typically there are two lines at any given time, one for each of the two active modules. When a module is completed, the fraction for that module goes to zero, e.g. the disappearing (black) line representing the fraction for *Clapping* at $t=146$ s in the top plot of [Figure 4](#). And when, a new module is introduced, a new line in a new colour emerges, e.g. the appearing (green) line representing *Whip It* at $t=146$ s in the top plot of [Figure 4](#).

4. Evolutionary dynamic model

Our investigations in the rehearsal described above were motivated by our analysis of a mathematical model that we derived using the replicator–mutator dynamics from evolutionary game theory (Bürger 1998). These dynamics describe a game played by a population of individuals, each of which has a fixed strategy, interacts randomly with others, and receives a reward, also called a payoff, which determines its success in the game. In terms of natural selection, payoff is interpreted as fitness and success as reproductive success: the higher the fitness associated with a strategy, the faster the strategy will reproduce (Nowak 2006).

Letting the interacting individuals in the game represent the interacting dancers, and the strategies represent the dance modules, the replicator–mutator dynamics provide a useful model of TMBO. One important reason is that the fitnesses, i.e. the rewards, in the replicator–mutator dynamics change over time; in fact, fitnesses depend explicitly on the changing distribution of the population over the different strategies, much like we expect artistic reward in TMBO to depend on the changing distribution of dancers over the active modules. Another reason is that the rates of reproduction of strategies depend not only on fitness but also on mutation, which in the equations is given by a random term; this allows us to use mutation to represent uncertainty and spontaneity associated with the dancers' choices. Further, the outcomes of the replicator–mutator dynamics range from mixed strategy solutions, where multiple strategies coexist, to pure strategy solutions, where a single strategy is adopted by the entire population. This parallels our central interest: the tension between how dancers exploit active modules and how they converge on a single module so that a new one can be explored.

The outputs of the replicator–mutator dynamics are the time-varying fractions of the population associated with the different strategies. This corresponds nicely to the dance, where we seek to understand how the fraction of the population of dancers performing each active module changes as a function of time (as in [Figure 4](#)). And, getting to the original motivation for our collaboration, the mathematical model offers a number of parameters that can be modulated to correspond to qualitative instructions, i.e. dials, that the dancers might receive either before or during a performance.

So we let the strategies in the replicator–mutator model represent the dance modules, and evolutionary time t represent the time t during the dance. We set the total number of strategies N equal to the limit on number of dance modules that can be performed at a time. For our study we considered the replicator–mutator dynamics with $N=2$. We interpreted a drop close to zero in the fraction of one strategy followed by a rise in that fraction as the completion of one dance module followed by the introduction of a new module. In this way, the replicator-mutator dynamics could be used to represent the dancers' progression through some or all of the available catalogue of modules, while respecting the rule limiting the number of modules allowed to be performed at a time.

In the following, we describe the model for the case of $N=2$. The more general model for $N \geq 2$ is provided in the *SI Appendix S2*. The replicator–mutator equations describe the changing fractions in a very large population, so we made the simplifying abstraction that the number of dancers in the group is very large. Then, the fraction of dancers $x_1(t)$ performing one of the active modules (call it module 1) at time t and the fraction of dancers $x_2(t)$ performing the other of the active modules (call it module 2) at time t can take any value in the interval from 0 to 1. Since $x_1(t) + x_2(t) = 1$ for all time t , given one fraction, e.g. $x_1(t)$, we can always find the other, as $x_2(t) = 1 - x_1(t)$.

The time rate of change of x_1 , denoted dx_1/dt , and the time rate of change of x_2 , denoted dx_2/dt , are given by the replicator-mutator dynamics:

$$\begin{aligned}\frac{dx_1}{dt} &= (x_1 f_1 q_{11} + x_2 f_2 q_{21}) - \phi x_1, \\ \frac{dx_2}{dt} &= (x_1 f_1 q_{12} + x_2 f_2 q_{22}) - \phi x_2.\end{aligned}\tag{1}$$

Here, $x_1(t)$ and $x_2(t)$ are written as x_1 and x_2 for brevity. And $x_1 f_1 q_{11}$ refers to the product (multiplication) of the three variables x_1 , f_1 , and q_{11} , and similarly for the other terms. The variables f_1 and f_2 are the fitnesses (payoffs or rewards) of module 1 and 2, respectively, and $\phi = f_1 x_1 + f_2 x_2$ is the average fitness over the two modules. q_{21} , respectively q_{12} , is the probability of mutating from module 2 to 1, respectively from module 1 to 2. q_{11} and q_{22} are the probabilities of not mutating. So the term in the first equation $x_1 f_1 q_{11} - \phi x_1$, which is equal to $x_1 \phi (q_{11} f_1 / \phi - 1)$, will contribute to growing x_1 if $q_{11} f_1 / \phi > 1$, i.e. if the relative fitness f_1 / ϕ of module 1 is sufficiently large, or shrinking x_1 if $q_{11} f_1 / \phi < 1$, i.e. if the relative fitness f_1 / ϕ of module 1 is sufficiently small. The other term $x_2 f_2 q_{21}$ will modulate the growth or shrinkage of x_1 by a term in the rate that is proportional to q_{21} , the probability of mutating from module 2 to 1.

Equation (1) was originally derived to describe evolutionary biology. However, we were motivated in part to generalize it for our purposes because it can also be derived as the limit of a stochastic error-prone imitation process, where agents imitate a strategy, e.g. strategy 1 with $x_1 > 0$, at a rate proportional to its relative fitness f_1 / ϕ and mutate (e.g. because of an error) to strategy 1 from alternative strategy 2 at a rate proportional to q_{21} (Traulsen, Claussen, and Hauert 2006).

In applications of the replicator–mutator dynamics, the fitnesses f_1 and f_2 are typically defined as *linear* functions of the fractions x_1 and x_2 , as in Nowak, Komarova, and Niyogi (2001). A linear fitness function implies that the sensitivity of fitness to changes in fractions is a constant. For example, if the fitness function is linear in x_1 , then if x_1 is doubled, the fitness doubles no matter if x_1 is a small, intermediate, or large number. We found that to provide a model rich enough to represent the dynamics of TMBO, we needed more subtle sensitivity of fitness to fractions. We thus generalized the fitness function so that it has low sensitivity to changes in x_1 and x_2 when they are very small (near 0) or very large (near 1) and high sensitivity to changes in x_1 and x_2 when x_1 and x_2 take intermediate values (near 1/2). The idea was to make the response highly sensitive only for intermediate values of x_1 and x_2 , when major qualitative transitions are prone to appear.

We defined the new nonlinear fitness functions using what's known as a Hill type function $\sigma_{\gamma,k}$ as follows (where γ and k are parameters):

$$\begin{aligned} f_1 &= b_{11}\sigma_{\gamma,k}(x_1) + b_{12}\sigma_{\gamma,k}(x_2), & \sigma_{\gamma,k}(x_1) &= \frac{\left(\frac{x_1}{1-x_1}\right)^\gamma}{k + \left(\frac{x_1}{1-x_1}\right)^\gamma}, \\ f_2 &= b_{21}\sigma_{\gamma,k}(x_1) + b_{22}\sigma_{\gamma,k}(x_2), & \sigma_{\gamma,k}(x_2) &= \frac{\left(\frac{x_2}{1-x_2}\right)^\gamma}{k + \left(\frac{x_2}{1-x_2}\right)^\gamma}. \end{aligned} \quad (2)$$

The coefficient $b_{11} \geq 0$ describes the strength of the dependence of f_1 on x_1 (b_{12} the dependence of f_1 on x_2 and likewise for b_{21} and b_{22}). For example, a large b_{12} could represent the relative ease in transitioning commitment from module 2 to module 1. The function $\sigma_{\gamma,k}(x_1)$ is 'sigmoidal' in x_1 (likewise $\sigma_{\gamma,k}(x_2)$ is sigmoidal in x_2), which means that it saturates at the value 1 for x_1 close to 1 and at the value 0 for x_1 close to 0. When $\gamma = k = 1$, it specializes to the linear function $\sigma_{1,1}(x_1) = x_1$ (and $\sigma_{1,1}(x_2) = x_2$). The parameters k and γ determine the shape of the sigmoidal function $\sigma_{\gamma,k}$. Notably, for smaller values of k , $\sigma_{\gamma,k}$ has a steeper slope (becomes more sensitive) and saturates at small or large x_1 (see *SI Appendix*, Figure S2).

The degree to which there is spontaneous or random switching between modules is governed by a probability μ called mutation strength, which takes a value in the interval from 0 to 1. We let

$$q_{11} = q_{22} = 1 - \mu, \quad q_{12} = q_{21} = \mu. \quad (3)$$

Then μ represents the probability of randomly switching from module 1 to 2 or 2 to 1, and $(1 - \mu)$ represents the probability of not switching.

Given a dynamic system defined by nonlinear equations, a bifurcation analysis can be used to reveal the sensitivity of the dynamical behaviour of the system to the value of a system parameter. A bifurcation refers to the change in number and stability of steady solutions of the dynamical equations as the parameter, called the bifurcation parameter, passes through a critical value, called the bifurcation point. Bifurcation analysis of the dynamics of equations (1)–(3), in the linear case $\gamma = k = 1$, show the solutions and their stability to be sensitive to the value of bifurcation parameter μ (Nowak, Komarova, and Niyogi 2001; Komarova and Levin 2010; Pais, Caicedo-Núñez, and Leonard 2012).

Consider the case of symmetric fitness coefficients in which $b_{11} = b_{22} = 1$ and $b_{12} = b_{21} = b$ with $0 < b < 1$. If μ is large enough, there is a single stable distributed equilibrium (mixed strategy solution) corresponding to $x_1 = x_2 = 0.5$, i.e. the population is distributed equally over the two modules. If μ is small enough, the distributed equilibrium is unstable, and there are two stable solutions, each corresponding to one strategy dominating the other in the population. In the limit as μ decreases to zero there is bi-stability of the two fully dominating module equilibria (pure strategy solutions) corresponding to $x_1 = 1, x_2 = 0$ and $x_1 = 0, x_2 = 1$ (Nowak, Komarova, and Niyogi 2001), i.e. the stable solutions correspond to the population fully committed to one or the other module. Thus, the magnitude of μ determines if the population is uniformly distributed over the modules or if a single module dominates.

As a model of TMBO, we interpret the distributed equilibrium, $x_1 = x_2 = 0.5$ as full exploitation, since it reflects juxtaposition of current modules with half the group performing module 1 and half the group performing module 2. We interpret the fully dominating module equilibria as the prerequisites for exploration, since they reflect the opportunity to introduce one or more new modules with the majority of the group performing one of the modules and a minority performing the other. The bifurcation analysis suggested that μ is a key control in the explore–exploit tradeoff. Since the explore–exploit tradeoff in the dance was observed to be dynamic, we augmented our model with dynamics for μ . That is, instead of letting μ be a fixed parameter, we defined feedback dynamics for μ that represent how the dancers might have been modulating something like μ based on their observations of how many dancers were performing each module (x_1 and x_2). Because dancers are trained to be highly physically aware, they will have good estimates, not only of the fractions x_1 and x_2 , but also of the presence or absence of a dominating module, i.e. whether the group is fully exploiting or in a position to explore a new module.

The dynamics of μ that we introduced into the model (Equation (4) below) allow flexibility in representing how the dancers react to the presence or absence of a dominating module, i.e. to the presence or absence of a large majority of the dancers performing one module. The intuition is that the model should allow for the dancers to cycle between the presence and absence of a dominating module, and in turn cycle between exploring and exploiting as artistic rewards for choices grow and decline over time. So we designed the dynamics such that if no module is dominating, μ decreases and drives the fractions towards a dominating module. Likewise, if one of the modules is dominating, μ increases and drives the fractions away from the dominating module solution. How sensitive the reaction is depends on the remaining model parameters, notably, the sensitivity parameter k in $\sigma_{\gamma,k}$. Thus, μ and k , or some combination of the two, might be used to translate the notion of switchiness from the rehearsal to the model, i.e. modulation of μ and k in the model might provide a useful representation of a switchiness dial for the dancers.

Let $K > 0$ be a time-scale parameter that regulates how fast μ changes relative to how fast x_1 and x_2 change. Let $0 < \alpha_1 < 0.5$ and $\alpha_2 = 1 - \alpha_1$ be thresholds. We define the time rate of change of μ , denoted $d\mu/dt$, as

$$\frac{d\mu}{dt} = K(x_1 - \alpha_1)(x_1 - \alpha_2)\mu(1 - \mu). \quad (4)$$

Since $x_2 = 1 - x_1$, the right side of Equation (4) can equivalently be written as $K(x_2 - \alpha_1)(x_2 - \alpha_2)\mu(1 - \mu)$. As intended these dynamics imply that if the fractions are close to the distributed solution where $\alpha_1 < x_1 < \alpha_2$ and $\alpha_1 < x_2 < \alpha_2$, then one of the terms in the product on the right side of Equation (4) will be negative so that μ will decrease and the fractions will move towards a dominating solution. Alternatively, if the fractions are close to a dominating solution where $0 < x_1 < \alpha_1$ and $\alpha_2 < x_2 < 1$ or $0 < x_2 < \alpha_1$ and $\alpha_2 < x_1 < 1$, then the product on the right side of Equation (4) will be positive so that μ will increase and the fractions will move towards a distributed solution. See the *SI Appendix S2* for a more general formulation of the dynamics of μ . An earlier version of these dynamics is described in Özcimder et al. (2016).

We note that the dynamics of μ depend not only on k but also on α_1 and α_2 . So α_1 and α_2 could be investigated as to their influence on the overall dynamics and thus also as

possible representatives of dials for the dancers. Indeed, in the model of Dey et al. (2018) we extend the model described here by introducing dynamics for α_1 and α_2 , which results in even more varied and rich dynamic behaviour.

5. Results from the studio

As discussed above, we quantified the explore–exploit dynamics of the group, in both human and modeled settings, in terms of the changing distribution of dancers over the modules. The changing distribution can be visualized by plotting the fraction of total number of dancers performing each of the current modules as a function of time. As seen in Figure 4, the top, middle, and bottom plots show the fractions x_1 and x_2 of the group of dancers in Run-through 1, 2 and 3, respectively, as a function of time.

During the 540 s of Run-through 1, as can be observed in the top plot of Figure 4, the dancers performed five of the nine available modules in the catalogue. At all times except for the beginning when the dancers were first entering and during the brief period at $t=318$ s, it can be observed that there were only two active modules. During the short interval at $t=318$ s, a third module (*Do Op*) was introduced a moment or two before the lingering dancer (Rhonda) had completed the second active module (*Lay Down and Get Up*).

The top plot of Figure 4 shows that one after another the dancers started with *Clapping* until, at $t=34$ s, there were six dancers in *Clapping* ($x_1 = 0.67$). One dancer then introduced *Lay Down and Get Up*, and others followed until at $t=57$ s all nine dancers were engaged in either *Clapping* or *Lay Down and Get Up*. From $t=70$ to 110 s, the nine dancers were evenly distributed over the two modules, with some switching of membership. Then, x_2 grew while x_1 shrunk, until at $t=135$ s *Clapping* was gone ($x_1 = 0$) and *Lay Down and Get Up* was fully dominating ($x_2 = 1$). At that moment one and then another dancer introduced the new module *Whip It* and two others joined so that at $t=170$ s, $x_3 = .44$ and $x_2 = .56$. This distributed state persisted for a while but then *Lay Down* dominated again ($x_2 = .89$). However, *Whip It* did not go away. Instead, a few seconds later, there was a return to the distributed state followed by a period in which *Whip It* dominated ($x_1 = .89$). At $t=318$ s *Do Op* was introduced and quickly dominated ($x_4 = 1$). A few second later *Drop and Roll* was introduced and there followed oscillations between a dominating module and a distributed state until at the end *Do Op* was fully dominating.

Over the 540 s shown in all three of run-throughs in Figure 4, the dancers can be observed to dynamically balance the explore–exploit tradeoff by oscillating as a group between a dominating module (prerequisite to exploration) and a distribution over two modules (exploitation). The frequency of the oscillations was irregular and the module that became dominant in each cycle was irregularly chosen. The order of introduction of modules was not necessarily the order of completion of modules. For example, in Run-through 1, *Lay Down* lasted as long as *Whip It* even though it was introduced well before and *Do Op*, introduced before *Drop and Roll*, outlasted *Drop and Roll*. In Run-through 2, *Clapping* outlasted *Bumper Body*, *Folk*, and *Romper Room*. In Run-through 3, *Box Drag* was the second module introduced and it was still active at the end of the 9-minute mark.

Although the explore–exploit tradeoff was not discussed with the dancers prior to the investigation in the studio, during the rehearsal, the dancers described feeling a distinct

tension between experimenting with current modules and moving on to introduce new modules. They also described finding new challenges and creative opportunities with the new rules, most especially the limit to two modules at a time. For example, dancers who resisted joining a popular module influenced the dynamics in new ways. This was very apparent when Rhonda resisted giving up on *Lay Down and Get Up* in Run-through 1. By exploiting the less-populated module, Rhonda actively blocked exploration.

Likewise, dancers who worked to recruit others to their module invented new variation. By recruiting to the more-populated module, these dancers actively advanced exploration. The snapshot in Figure 2(e) shows eight dancers in a circle around Rhonda, on the floor, clearly urging her through their physical orientation and facial expressions to join them in *Whip It*. In Run-through 3, Chris went up to several dancers performing *Folk* and folded them into *Box and Drag* until no one remained in *Folk*. He then very quickly introduced *Bravo* before anyone could go back to *Folk*. Later Cori created a moment evocative of Chris' strategy by folding dancers who were performing *Bravo*. However, Cori did not appear to be insistent and it read as playful and humorous.

The constraints imposed by the modified rules and heightened tension were seen to motivate new dramatic moments. And because the dancers were allowed to switch between active modules, rather than being required to follow the order that modules were introduced, they reported that they could hold on to an idea and find a time to switch back to a module to pursue their idea when they saw the chance. As a result of the experience of the investigation in the studio, the new rules became an ongoing part of the practice and the language of explore-exploit became a baseline for development of the piece. In performance, even when the limit of two modules was lifted and the other modifications removed, the dancers created moments inspired by these rules.

Run-throughs 2 and 3 were also used to investigate the possibility of giving directions, like a design dial, to the dancers that would be analogous to manipulating design parameters in the model (like μ and k) to modify the collective behaviour. This was investigated by asking the dancers to modify their switching tendency. The results suggested that switching tendency could be modified through instruction and would shape the dynamics in interesting ways. For example, the different instructions on switching tendency were reflected in the different total numbers of modules introduced in the first 480 s of the improvisations: five modules in Run-through 1, six in Run-through 2, and four in Run-through 3. The higher number of modules is consistent with impulsiveness, i.e. a dial up in switching tendency (the instruction in Run-through 2) and the lower number with resistance, i.e. a dial down in switching tendency (the instruction in Run-through 3). A similar effect can be observed in Figure 5, which shows all of the times at which x_1 (and therefore x_2) switched from increasing to decreasing or vice versa for each of the three run-throughs. These switches were most frequent and numerous for Run-through 2, when switching tendency was high, which yielded a total of 29 switches. They were least frequent and numerous for Run-through 3, when switching tendency was low, which yielded a total of 20 switches. We note, however, that this analysis might even underestimate the effect of the switching tendency instruction, since it does not account for switches by dancers that did not change the overall fractions. For example, in Run-through 2, dancers did a lot of switching between modules, e.g. between *Clapping* and *Folk*. But often one dancer in *Folk* would return to *Clapping* just as another dancer in *Clapping* took their place in *Clapping*, thus leaving the overall fractions unchanged.

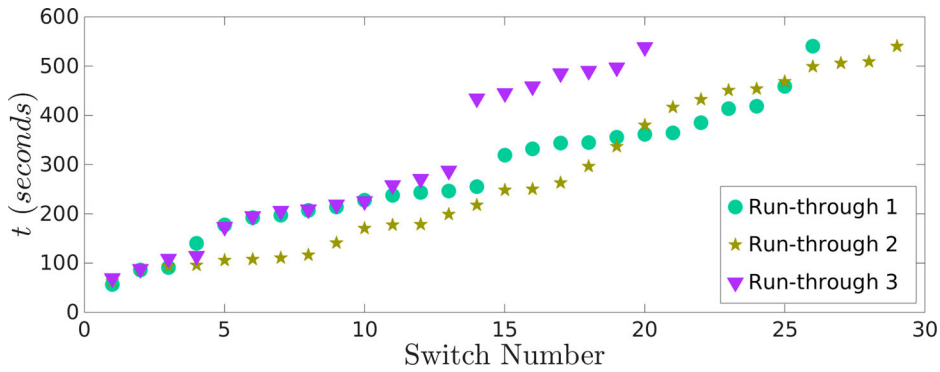


Figure 5. Switching times for Run-throughs 1, 2, and 3. Each switching time refers to a time when the value of x_1 (and therefore x_2) changed from increasing to decreasing or vice versa.

6. Results from the model

A bifurcation analysis of the model was used to investigate the roles of μ and k in the modeled behaviour of the group. The analysis influenced, and was influenced by, the investigation in rehearsal. To visualize the results of the analysis, we used a bifurcation diagram, which shows how steady solutions of the model dynamics change for different values of the bifurcation parameter μ , which in turn depends on k . We considered the symmetric case in which $b_{11} = b_{22} = 1$, $b_{12} = b_{21} = b$. The steady solution is represented by either x_1 or x_2 , and, since $x_1 + x_2 = 1$, if one steady solution is x_i , then the other steady solution is $1 - x_i$. So the bifurcation diagram is a plot of steady solutions x_i (think of it as x_1) on the vertical axis versus bifurcation parameter μ on the horizontal axis.

Six distinct bifurcation diagrams for equations (1)–(3) with $b = 0.04$, are plotted in the top panel of Figure 6. These plots show how the steady solutions x_1 (and $x_2 = 1 - x_1$) and their stability vary over a range of values of μ for six different values of k (see *SI Appendix S4* for a proof). Each of the six bifurcation diagrams is plotted in a different colour as indicated by the key. Solid lines indicate solutions that are stable and dashed lines indicate solutions that are unstable. If the solution is stable, then it remains steady even in the case of a small change in conditions. If the solution is unstable, however, even the smallest change in conditions will drive the dynamics away from the unstable solution and towards a stable solution.

For every value of μ and k , there is a steady solution at the distributed solution $x_i = 0.5$, corresponding to half the population in module 1 and half in module 2. This means that all six bifurcation diagrams have horizontal lines at $x_i = 0.5$. To avoid overlapping lines from the different diagrams, we only show the solution $x_i = 0.5$, for each of the six bifurcation diagrams, along the line near its intersection with the curved part of the diagram. This point of intersection is known as the pitchfork bifurcation point. We denote by $\mu_c(k)$ the value of μ at the pitchfork bifurcation point; we write μ_c as a function of k because the location of the pitchfork bifurcation point depends on the value of k . In Figure 6, we identify μ_c for $k = 0.05$ in orange and μ_c for $k = 0.4$ in darkest blue. In each bifurcation diagram, the steady solution $x_i = 0.5$ is stable for $\mu > \mu_c$ (solid line) and unstable for $\mu < \mu_c$ (dashed line). Thus, in Figure 6 every solid line on $x_i = 0.5$ continues to the right (as μ increases) and every dashed line on $x_i = 0.5$ continues to the left (as μ decreases).




Figure 6. Evolutionary model with $N=2$, $b=0.04$, $\gamma = 2.5$. (Top) Bifurcation diagrams of Equations (1)–(3) showing steady solutions x_i versus μ for six different values of k in the nonlinear fitness function. Solid lines are stable solutions and dashed lines are unstable solutions. (Middle) Bifurcation diagram from top panel for $k=0.05$ with green and purple curves, each showing path of solution x_i , for a different initial condition, and for increasing and then decreasing μ . (Bottom) x_i as a function of time for Equations (1)–(4) with $K=0.05$, $\alpha_1 = 0.25$, $\alpha_2 = 0.75$, and noise variance $S(\cdot)$ a Gaussian with mean 0.5 and standard deviation 0.05.

When $k = 0.4$ the bifurcation diagram (farthest to the right in darkest blue) reveals a ‘supercritical’ pitchfork bifurcation, represented by the c-shaped curve that opens to the left, at bifurcation point $\mu = \mu_c(0.4) \approx 0.28$. The diagram shows that for $\mu > \mu_c$ the distributed solution $x_i = 0.5$ is the only (stable) solution. For $\mu < \mu_c$ the distributed solution

is unstable and the two (symmetric) dominant solutions are stable: $x_1 > 0.5$ ($x_2 < 0.5$) and $x_1 < 0.5$ ($x_2 > 0.5$). The magnitude of the stable dominant solutions increases smoothly as μ decreases. Thus, were μ to slowly decrease (increase) past the bifurcation point, the result would be a slow change from the distributed solution to a dominant solution (vice versa).

The results from the studio in Figures 4 and 5 show instead relatively rapid switching in x_i , corresponding to oscillations in exploration versus exploitation. This suggests that the model with $k \geq 0.4$, which includes the linear fitness model from the literature, is insufficient to fully capture the observed behaviour. However, in our model with nonlinear fitness function, rapid switching does follow for slowly varying μ if k is sufficiently small, much like the appearance of oscillatory behaviour in relaxation oscillators (see *SI Appendix S5*). This can be understood from the bifurcation diagrams in the top plot of Figure 6 where $k < 0.4$.

For example, when $k = 0.05$ the bifurcation diagram (second from the right in orange) reveals a ‘subcritical’ pitchfork bifurcation, distinguished by the c-shaped curve that opens to the right, at bifurcation point $\mu_c = \mu_c(0.05) \approx 0.09$. As in the supercritical case, when $\mu < \mu_c$ the distributed solution is unstable and the two (symmetric) dominant solutions are stable. However, the situation is a little more complicated when $\mu > \mu_c$. Here, before μ gets very large, there are five steady solutions, three stable (solid) and two unstable (dashed). It is this more complicated diagram that leads to the relatively rapid switching in the fractions x_1 and x_2 as μ changes.

We illustrate this with the clockwise (green) and counterclockwise (purple) curves in the middle plot of Figure 6, which are model simulations plotted on top of the bifurcation diagram in the case $k = 0.05$. The simulated curves show how the solution x_1 changes for two different initial conditions as μ is slowly increased and then slowly decreased past the subcritical bifurcation point. The initial conditions for the green curve were approximately $x_1 = 0.8$ and $\mu = 0.02$, and x_1 can be seen first to grow closer to 1 then to shrink to 0.5 and then to grow back towards 1. The initial conditions for the purple curve were approximately $x_1 = 0.1$ and $\mu = 0.08$, and x_1 can be seen first to shrink closer to 0 then to grow to 0.5 and then to shrink back towards 0. Each loop represents a cycle between exploration and exploitation in the modeled dynamics. In the green curve, it is module 1 that the population converges on before exploring a new module, whereas in the purple curve it is module 2.

The corresponding oscillatory behaviour of x_1 as a function of time t is represented in the bottom plot of Figure 6: one circuit around the green periodic solution in the middle plot corresponds to one up-down oscillation in the bottom plot and one circuit around the purple periodic solution in the middle plot corresponds to one down-up oscillation in the bottom plot.

The plot of x_1 versus t in the bottom panel of Figure 6 is a simulation of Equations (1)–(4) for $N = 2$, $b = 0.04$, $k = 0.05$, $\gamma = 2.5$, $K = 0.05$, $\alpha_1 = 0.25$, $\alpha_2 = 0.75$. Here μ is not independently driven up and down, but rather it follows its own dynamics as prescribed by Equation (4). That is, Equations (1)–(4) are the dynamics of the distribution of the population between the two active modules with small k and feedback dynamics for μ corresponding to more switching when one module is dominant and less switching when no module is dominant. These dynamics yield relatively rapid oscillations as observed in the rehearsal.

In this simulation, white noise is added to Equation (1) with variance $S(x_i)$ that depends on x_i to represent uncertainties (like that observed in Figure 4 at $t = 318$ s). $S(\cdot)$ is designed to be symmetric about $x = 0.5$, vanishing on the boundaries, i.e. $S(0) = S(1) = 0$, and greatest at $x_i = 0.5$, i.e. when it is more difficult for dancers to judge distribution across modules. The noisy response around $x_i = 0.5$ can be observed in the simulation of Figure 6; it leads to switching between the two different periodic solutions (green and purple in the middle plot).

We note further, from the top of Figure 6, that as k increases the magnitude of oscillations decreases, and oscillations exist for higher values, and a smaller range, of μ . The insight here is that a small change in k can shape the oscillations in the explore-exploit dynamics of the model.

To make a qualitative comparison of the model output with observations from the studio (Figure 4), we plot in Figure 7 a model simulation of x_1 and x_2 as a function of time t , discretized for a population of nine. Only two modules are allowed at a time and once x_1 or x_2 goes to zero, the corresponding module is replaced by a new module indicated by a new colour. The simulation uses the same parameters as in Figure 6, except that k is varied over time, as shown in the bottom panel of Figure 7, to represent variability in sensitivity over time.

The simulation in Figure 7 illustrates the influence of μ and k on the explore-exploit tradeoff; it can be seen that changes in these parameters affected the model dynamics much the same way as the change in instruction in Run-throughs 2 and 3 (for more and less tendency to switch) affected the collective behaviour as seen in the middle and




Figure 7. Model simulation results for comparison with results in rehearsal plotted in Figure 4. (Top) The fraction x_i of a population of nine as a function of time t for a simulation of the model with $N=2$, $b=0.04$, $\gamma = 2.5$, $K=0.05$, $\alpha_1 = 0.25$, $\alpha_2 = 0.75$, $S(\cdot)$ a Gaussian with mean 0.5 and standard deviation 0.05, and time-varying k . When x_i goes to zero, a new module is introduced. Different modules are distinguished by different colours. (Bottom) Parameter k varies with time t .

bottom plots of [Figure 4](#). For example, during the period from $t = 100$ to 140 s in the simulation of [Figure 7](#), k was dramatically increased resulting in a pause in oscillations in x_1 and x_2 . We observe in [Figure 7](#) features in common with all of the run-throughs in [Figures 4](#): oscillations with irregular frequency, irregular choice of dominating module in each cycle, and different ordering of introduced versus completed modules.

7. Conclusions and future directions

Until we began our experiments together, TMBO existed as a dance without music, imposing the general rule of staying within three modules of one another and not allowing skipping or switching between modules. Following the underlying principle that the dance should be different every time and the dancers should be looking for something different in the moment, we employed new approaches to examine how the rules and constraints could be modified and designed to support creativity and bring forth new experiences. By tuning our collective attention to understanding the explore versus exploit tension, where spontaneity and risk-taking could be artistically rewarding, we widened the dramatic possibilities of the score and heightened the performers' knowledge of how they prefer to engage within the group, which then created new possibilities for both openness and control. We permanently adapted our rules to include sections where external mechanisms signal a time to adapt responsive behaviours, i.e. to turn the 'dials' in one direction or another. In the year following our joint work, Dan Trueman completed the composition of forty-four music modules in collaboration with the music ensembles So Percussion and Mobius Percussion, and we produced a fully realized version of TMBO with 15 dancers and 12 percussionists at New York Live Arts in March 2016. TMBO has since been reprised on several occasions, including at The Scotia Festival of Music, Nova Scotia, Canada, in June 2017 (for a video of the performance, see [Lazier 2017](#)).

In becoming an integral part of the development of the dance, our art and science investigation of TMBO produced new perspectives on social decision-making dynamics and opportunities for design in a real-time, collaborative, art-making context. By integrating studies with the TMBO dancers in rehearsal and analysis of a representative mathematical model, we focused in on a simply parameterized mechanism to help describe and shape how dancers in TMBO address an explore–exploit tension that derives from competing artistic goals and constraints imposed by the choreographic rules.

In most contexts studied in the explore–exploit literature, rewards associated with decision-making options do not change with time. Decision-makers explore and exploit until they find the best option, and then they stop exploring. In TMBO, however, the rewards continue to change and so the performers continue to explore. An active module that is appealing or provocative at one moment will eventually become less so with time or with changing circumstances. A performer may first exploit the active module and then switch to exploring by introducing or joining in on a new module. The new module may in turn lead others to exploit. By switching between exploring and exploiting, the dancers keep fresh the sequencing, timing, counterpoint, relationship, and variation of modules. The dancers' switching behaviour was observed in the oscillatory dynamics of the fractions x_1 and x_2 of the total number of dancers committed to each of the two active modules ([Figure 4](#)).

The oscillatory dynamics in x_1 and x_2 were likewise captured with a mathematical model, generalized from the replicator-mutator equations of evolutionary biology to include a nonlinear fitness function with sensitivity parameter k and self-contained feedback dynamics for mutation rate μ . For small enough k , the model reveals a symmetric subcritical pitchfork bifurcation, which provides a mechanism for rapid switching between exploration and exploitation as μ evolves. Modulation of the parameters μ and k in the model appeared to account well for the subtlety and variability of the dancers' innovative impulses in reactions to what the others were doing, thus providing a possible representation of instruction dials that could be used to signal rule changes to the dancers. For example, changes in k , which in turn changed the dynamics of μ , were shown to affect the frequency of the explore–exploit oscillations (Figure 7), much like what was observed in rehearsal (Figure 4) when dancers were asked to modulate their switchiness tendency. In this case, a simple implementation of such a dial would be to define one beat of a drum as an instruction to the dancers to increase their switchiness tendency and two beats of the drum as an instruction to the dancers to decrease their switchiness tendency.

The model can be used further as an off-line choreographic tool to examine the predicted consequences on collective explore–exploit dynamics of other kinds of design modifications. For example, study of the influence of asymmetries in the relative ease in transitioning from one module to another can be made using the parameters b_{ij} (Komarova and Levin 2010). And study of the influence of thresholds in the shifting dynamics of mutation rate can be made using the parameters α_1 and α_2 (Dey et al. 2018). The case in which the limit on the number of current modules is greater than two can be studied using the replicator–mutator dynamics with $N > 2$ (Pais, Caicedo-Núñez, and Leonard 2012).

Our investigation, insights, and mathematical model may generalize to other contexts beyond structured improvisational dance, in the same way that the replicator–mutator dynamics have been applied in various contexts. For example, the simply parameterized but rich family of collective dynamics may provide a useful framework for examining mechanisms of social decision-making in settings where expert performers together face an explore–exploit tension in which rewards change with time or with changing circumstances, such as in other art forms, collaborative design, or team sports. The generalized model advances a related goal, which is to draw inspiration from the dynamics of natural systems and translate the mechanisms uncovered, using mathematics, to models that can be used for design (Gray et al. 2018). In all these ways, our collaboration has proved to be a generative meeting of art and science.

Acknowledgments

The authors thank everyone who contributed to making TMBO, most especially the dancers. K. O. and N. E. L. designed the research. K. O., R. L., D. T., and N. E. L. performed the research. K. O., B. D., A. F., and N.E.L. contributed new analytic tools. K. O., B. D., A. F., and N. E. L. analysed the data. N. E. L. wrote the paper with input from all co-authors.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This research was supported in part by a grant for Artists and Scientists or Engineers from the Office of the Dean for Research at Princeton University.

Notes on contributors

Kayhan Özcimder received his bachelor's degree (2007) in Mechanical Engineering from Gazi University, Ankara, Turkey and his PhD (2015) and MSc degree (2012) in Mechanical Engineering from Boston University. Since February 2015 Kayhan has been with the Department of Mechanical and Aerospace Engineering and Princeton Neuroscience Institute at Princeton University. His primary research interests are bio-inspired control, animal and human collective behavior and communication-through-motion in multi-agent systems.

Biswadip Dey received his bachelor's degree in Electrical Engineering from Jadavpur University, India, in 2006, master's degree in Systems & Control Engineering from the Indian Institute of Technology Bombay, India, in 2008, and PhD in Electrical Engineering from the University of Maryland, College Park, USA, in 2015, respectively. Since March 2015, he has been with the Department of Mechanical and Aerospace Engineering at Princeton University, where he is an Associate Research Scholar at present. His primary research interests focus on nonlinear, geometric and optimal control, and their applications in collective robotics and autonomous systems.

Alessio Franci received his MS in Theoretical Physics from the University of Pisa in 2008 and his PhD in Physics and Control Theory from the University of Paris Sud 11 in 2012. Since April 2015 he is Associate Professor of Biomathematics at UNAM. His research interests span different fields but the focus is on the regulation principles of biological behaviors, in particular, rhythmic and excitable neuronal behaviors (from single neurons to cognition) and collective decision making. Thanks to local and international collaborations, his research has a marked interdisciplinary inclination that has recently started to open up to the Arts.

Rebecca Lazier choreographs and directs a project-based group of dancers in New York and is a Senior Lecturer in dance at Princeton University. A Julliard alum, she has produced her work in Hartford, Los Angeles, Istanbul, and New York. A film of her piece *Coming Together/Attica*, which premiered in June 2013, was featured in Venice, Italy as part of the Architecture Biennale. In New York Lazier's work has been presented at venues including La MaMaMoves! Festival, DanceSpace Project, The Kitchen, the Guggenheim Museum, 92nd Street Y, Joyce SoHo, and Movement Research at the Judson Church. The company has toured to a variety of locales from Martha's Vineyard to Los Angeles, Jacob's Pillow to New Orleans, and Nova Scotia to Turkey.

Dan Trueman is a composer, fiddler, and electronic musician. Dan's collaborators include: So Percussion, fiddler Caoimhín Ó Raghallaigh, guitarist Monica Muga, the JACK Quartet, singer Iarla Ó Lionáird, bitKlavierist Adam Sliwinski, choreographer Rebecca Lazier, scientist Naomi Leonard, the Princeton Laptop Orchestra, Eighth Blackbird, and others. His work has been recognized by the Guggenheim Foundation, the Barlow Endowment, the Bessies, the Fulbright Commission, the American Composers Forum, the American Council of Learned Societies, Meet the Composer, among others. He is Professor of Music and Director of the Princeton Sound Kitchen at Princeton University, where he teaches counterpoint, electronic music, and composition.

Naomi Ehrlich Leonard works on dynamics and control for multi-agent systems, both technological and natural. She led a multidisciplinary project that culminated in a first-of-its-kind autonomous ocean observing system, featuring a coordinated network of underwater robots. With biologists, she studies collective animal behavior and human decision-making. With artists, she explores dance, music and feedback dynamics. Recognitions include a MacArthur Fellowship and membership in the American Academy of Arts and Sciences. She is Edwin S. Wilsey Professor of Mechanical and Aerospace Engineering, associated faculty in Applied and Computational Mathematics, and Director of the Council on Science and Technology at Princeton University.

ORCID

Alessio Franci  <http://orcid.org/0000-0002-3911-625X>

Naomi Ehrich Leonard  <http://orcid.org/0000-0002-5328-3871>

References

- Arrow, K. 1951. *Social Choice and Individual Values*. New York: John Wiley & Sons.
- Auer, P., N. Cesa-Bianchi, and P. Fischer. 2002. "Finite-time Analysis of the Multiarmed Bandit Problem." *Machine Learning* 47 (2): 235–256.
- Bürger, Reinhard. 1998. "Mathematical Properties of Mutation-Selection Models." *Genetica* 102–103 (0): 279–298.
- Carl, R. 2009. *Terry Riley's In C*. Oxford: Oxford University Press.
- Clemente, K. 1990. "Playing with Performance: The Element of the Game in Experimental Dance and Theater." *The Journal of Popular Culture* XXIV: 1–10.
- Cohen, J. D., S. M. McClure, and A. J. Yu. 2007. "Should I Stay Or Should I Go? How the Human Brain Manages the Trade-Off Between Exploitation and Exploration." *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 362 (1481): 933–942. <http://rspb.royalsocietypublishing.org/content/362/1481/933>.
- Conradt, L., and T. J. Roper. 2003. "Group Decision-Making in Animals." *Nature* 421: 155–158.
- Couzin, I. D., J. Krause, N. R. Franks, and S. A. Levin. 2005. "Effective Leadership and Decision-Making in Animal Groups on the Move." *Nature* 433: 513–516.
- Dey, B., A. Franci, K. Özçimder, and N. E. Leonard. 2018. "Feedback Controlled Bifurcation of Evolutionary Dynamics with Generalized Fitness." Proceedings of the American Control Conference (ACC), Milwaukee, WI, June 27–29, 6049–6054.
- Dyer, J. R. G., A. Johansson, D. Helbing, I. D. Couzin, and J. Krause. 2009. "Leadership, Consensus Decision Making and Collective Behaviour in Humans." *Philosophical Transactions of the Royal Society B: Biological Sciences* 364 (1518): 781–789.
- Fehr, E., and C. F. Camerer. 2007. "Social Neuroeconomics: The Neural Circuitry of Social Preferences." *Trends in Cognitive Sciences* 11: 419–427.
- Gray, R., A. Franci, V. Srivastava, and N. E. Leonard. 2018. "Multiagent Decision-Making Dynamics Inspired by Honeybees." *IEEE Transactions on Control of Network Systems* 5 (2): 793–806.
- Greening, B. R., Jr, N. Pinter-Wollman, and N. H. Fefferman. 2015. "Higher-Order Interactions: Understanding the Knowledge Capacity of Social Groups Using Simplicial Sets." *Current Zoology* 61: 114–127.
- Guimerà, R., B. Uzzi, J. Spiro, and L. A. Nunes Amaral. 2005. "Team Assembly Mechanisms Determine Collaboration Network Structure and Team Performance." *Science* 308: 697–702.
- Kalathil, D., N. Nayyar, and R. Jain. 2014. "Decentralized Learning for Multiplayer Multiarmed Bandits." *IEEE Transactions on Information Theory* 60 (4): 2331–2345.
- Kaufmann, E., O. Cappé, and A. Garivier. 2012. "On Bayesian Upper Confidence Bounds for Bandit Problems." International Conference on Artificial Intelligence and Statistics, La Palma, Canary Islands, Spain, April, 592–600.
- Kirsh, D. 2011. "Creative Cognition in Choreography." Proceedings of the 2nd International Conference on Computational Creativity, Mexico City, Mexico, April 27–29, 1–6.
- Kolla, R. V., K. Jagannathan, and A. Gopalan. 2016. "Stochastic Bandits on a Social Network: Collaborative Learning with Local Information Sharing." *CoRR*. <http://arxiv.org/abs/1602.08886>.
- Komarova, N. L., and S. A. Levin. 2010. "Eavesdropping and Language Dynamics." *Journal of Theoretical Biology* 264 (1): 104–118.
- Krebs, J. R., A. Kacelnik, and P. Taylor. 1978. "Test of Optimal Sampling by Foraging Great Tits." *Nature* 275 (5675): 27–31.
- Lai, T. L., and H. Robbins. 1985. "Asymptotically Efficient Adaptive Allocation Rules." *Advances in Applied Mathematics* 6 (1): 4–22.

- Landgren, P., V. Srivastava, and N. E. Leonard. 2016. "Distributed Cooperative Decision-Making in Multiarmed Bandits: Frequentist and Bayesian Algorithms." *Proceedings of the IEEE Conference on Decision and Control (CDC)*, Las Vegas, NV, December 12–14, 167–172.
- Lazier, R. 2017. *There Might Be Others at The Scotia Music Festival*. <https://vimeo.com/236973952>.
- Lazer, D., and A. Friedman. 2007. "The Network Structure of Exploration and Exploitation." *Administrative Science Quarterly* 52: 667–694.
- Lazier, R., D. Trueman, M. Baranova, and L. DeSilva-Johnson. 2016. *There Might Be Others*. Operating System.
- Lee, D. 2008. "Game Theory and Neural Basis of Social Decision Making." *Nature Neuroscience* 11 (4): 404–409.
- Leonard, N. E., T. Shen, B. Nabet, L. Scardovi, I. D. Couzin, and S. A. Levin. 2012. "Decision Versus Compromise for Animal Groups in Motion." *Proceedings of the National Academy of Sciences* 109 (1): 227–232.
- Leonard, N. E., G. F. Young, K. Hochgraf, D. T. Swain, A. Trippe, W. Chen, K. Fitch, and S. Marshall. 2014. "In the Dance Studio: An Art and Engineering Exploration of Human Flocking." In *Controls and Art*, edited by Amy LaViers and Magnus Egerstedt, 51–72. Springer. <https://arxiv.org/abs/1808.07842>.
- Mason, W., and D. J. Watts. 2009. "Collaborative Learning in Networks." *Proceedings of the National Academy of Sciences* 109 (3): 764–769.
- Meunier, H., J. B. Leca, J. L. Deneubourg, and O. Petit. 2006. "Group Movement Decisions in Capuchin Monkeys." *Behaviour* 143: 1511–1527.
- Nowak, M. 2006. *Evolutionary Dynamics: Exploring the Equations of Life*. Cambridge, MA: Harvard University Press.
- Nowak, M., N. L. Komarova, and P. Niyogi. 2001. "Evolution of Universal Grammar." *Science* 291: 114–118.
- Olfati-Saber, R., J. A. Fax, and R. M. Murray. 2007. "Consensus and Cooperation in Networked Multi-Agent Systems." *Proceedings of the IEEE* 95: 215–233.
- Özçimder, K., B. Dey, R. J. Lazier, D. Trueman, and N. E. Leonard. 2016. "Investigating Group Behavior in Dance: An Evolutionary Dynamics Approach." *Proceedings of the American Control Conference (ACC)*, Montreal, CA, July 6–8, 6465–6470.
- Pais, D., C. H. Caicedo-Núñez, and N. E. Leonard. 2012. "Hopf Bifurcations and Limit Cycles in Evolutionary Network Dynamics." *SIAM Journal on Applied Dynamical Systems* 11 (4): 1754–1784.
- Reed, S. Alexander. 2011. "In C on Its Own Terms: A Statistical and Historical View." *Perspectives of New Music* 49 (1): 47–78.
- Riley, T. 1964. *In C: Performing directions*. Tucson, AZ: Celestial Harmonies.
- Saari, D. G. 2001. *Decisions and Elections: Explaining the Unexpected*. Cambridge: Cambridge University Press.
- Sanfey, A. G. 2007. "Social Decision-making: Insights from Game Theory and Neuroscience." *Science* 318 (5850): 598–602.
- Seeley, T. D., and S. C. Buhrman. 1999. "Group Decision Making in Swarms of Honey Bees." *Behavioral Ecology and Sociobiology* 45: 19–31.
- Sen, A. K. 1970. *Collective Choice and Social Welfare*. San Francisco, CA: Holden-Day.
- Traulsen, A., J. C. Claussen, and C. Hauert. 2006. "Coevolutionary Dynamics in Large, But Finite Populations." *Physical Review E* 74: 011901.
- Villar, S. S., J. Bowden, and J. Watson. 2015. "Multi-armed Bandit Models for the Optimal Design of Clinical Trials: Benefits and Challenges." *Statistical Science* 30 (2): 199–215.

Social decision-making driven by artistic explore-exploit tension

Supporting Information

K. Özcimder, B. Dey, A. Franci, R. Lazier, D. Trueman, and N. E. Leonard

S.1 Synchronized video frames



Figure S1: Snapshots at the same time from each of the two synchronized videos of the dancers during Run-through 1.

S.2 General model for $N \geq 2$

In the main text we specialize the model to $N = 2$. Here we present here the general model for the case that the limit on number of modules at a time is $N \geq 2$. We make the simplifying abstraction that the number of dancers in the group is very large so that we can represent the fraction of dancers $x_i(t)$ committed at time t to strategy i as a number in the interval $[0, 1]$ for $i = 1, \dots, N$. By definition $\sum_{i=1}^N x_i(t) = 1$ for all time t . The time rate of change of x_i is given by the replicator-mutator dynamics:

$$\frac{dx_i}{dt} = \sum_{j=1}^N x_j f_j q_{ji} - \phi x_i, \quad (\text{S1})$$

where f_i is the fitness of strategy i , q_{ji} is the probability of mutating from strategy j to i , and $\phi = \sum_{i=1}^N f_i x_i$ is the average fitness over the N strategies.

We define a new nonlinear fitness function f_i using a Hill type function $\sigma_{\gamma,k}(x_j)$ as follows:

$$f_i = \sum_{j=1}^N b_{ij} \sigma_{\gamma,k}(x_j), \quad \sigma_{\gamma,k}(x_j) = \frac{\left(\frac{x_j}{1-x_j}\right)^\gamma}{k + \left(\frac{x_j}{1-x_j}\right)^\gamma}. \quad (\text{S2})$$

The coefficients $b_{ij} \geq 0$ describe the strength of the dependence of f_i on x_j . The function $\sigma_{\gamma,k}(x_j)$ is sigmoidal in x_j and specializes to the linear function $\sigma_{1,1}(x_j) = x_j$ for $\gamma = k = 1$.

Spontaneous switching between modules is governed by a probability $\mu \in [0, 1]$ called mutation strength. Following [1], we let

$$q_{ii} = 1 - \mu \text{ for all } i, \quad q_{ij} = \frac{\mu b_{ij}}{\sum_{l \neq i} b_{il}} \text{ for } i \neq j. \quad (\text{S3})$$

Then the probability that dancers spontaneously switch from module i to any other module is represented by $\sum_{j=1, j \neq i}^N q_{ij} = \mu$.

In the linear, symmetric case in which $b_{ii} = 1$ for all i and $b_{ij} = b \in (0, 1)$ for $i \neq j$, if μ is large enough there is a single stable distributed equilibrium corresponding to $x_i = 1/N$ for all i , i.e., the population is distributed equally over all strategies. If μ is small enough the distributed equilibrium is unstable, and the stable solutions correspond to one dominant strategy. In the limit as μ goes to zero there is multi-stability of a fully dominating strategy equilibria corresponding to $x_i = 1, x_j = 0, j \neq i$, i.e., the population is fully committed to one strategy.

Let $K > 0$ be a time-scale parameter and α a real-valued function of x_i . To define dynamics for μ , we introduce the scalar time-dependent variables $w_i(t)$ for $i = 1, \dots, N$ and let $\mu(t)$ depend on the $w_i(t)$. The time rate of change of w_i is

$$\frac{dw_i}{dt} = K\alpha(x_i)w_i(1 - w_i), \quad w_i(0) \in [0, 1].$$

So $w_i(t) \in [0, 1]$ for all $t \geq 0$ and equilibria are $w_i = 0$ and $w_i = 1$. We can design the function α using thresholds on the x_i to model how the dancers react to the presence or absence of a dominating module, i.e., to the presence or absence of a large majority of the dancers performing one module. For example, reinforcement of a dominant module can be modeled with $\mu = \min_i w_i$ and $\alpha(x_i) = \bar{\alpha} - x_i$ for $\bar{\alpha} \in (0.5, 1)$. Then for any i such that $x_i > \bar{\alpha}$, w_i and thus μ will decrease, driving the state dynamics toward a dominating module equilibrium. If $x_i \leq \bar{\alpha}$ for all i , then all the w_i and thus μ will increase, driving the state dynamics toward the distributed solution.

Alternatively, we can design the function α to model the opposite reaction of the dancers to the presence or absence of a dominating module. We do this in the main text for $N = 2$ in Eq. (4).

S.3 Nonlinear fitness function

The nonlinear fitness function given by Eq. (1) in the main text depends on the sigmoidal function $\sigma_{\gamma, k}(x_j)$. The parameters γ and k determine the shape of $\sigma_{\gamma, k}(x_j)$. We illustrate the influence of k on $\sigma_{\gamma, k}$ in Fig. S2.




Figure S2: The sigmoidal function $\sigma_{\gamma, k}(x_j)$ for different values of k and $\gamma = 2.5$.

S.4 Proof of symmetric quintic bifurcation

The proof of the supercritical pitchfork bifurcation is well known for the replicator mutator equations Eqs. (1)-(3) from the main text with $N = 2$ in the case of linear fitness function corresponding to $k = 1$. The transition from the supercritical to the subcritical pitchfork illustrated in Fig. 6 is through a symmetric quintic pitchfork bifurcation with normal form $-x^5 + \mu x$ [2, Section VI.5 and Figure VI.6.1 top]. We prove the existence of a

symmetric quintic pitchfork bifurcation for the replicator mutator equations Eqs. (1)-(3) from the main text with $N = 2$ in the case of the new generalized nonlinear fitness function corresponding to small k . In the case of $N = 2$, $b_{11} = b_{22} = 1$, and $b_{12} = b_{21} = b \in (0, 1)$, Eqs. (1)-(3) specialize to

$$\begin{aligned}\dot{x}_1 &= g(x_1, \mu, k, \gamma, b) \\ &:= x_1(\sigma_{\gamma,k}(x_1) + b\sigma_{\gamma,k}(1-x_1))(1-\mu) + (1-x_1)(b\sigma_{\gamma,k}(x_1) + \sigma_{\gamma,k}(1-x_1))\mu - \phi x_1,\end{aligned}\quad (\text{S4})$$

where

$$f_i = \sum_{j=1}^N b_{ij} \sigma_{\gamma,k}(x_j),$$

and

$$\sigma_{\gamma,k}(x_j) = \frac{\left(\frac{x_j}{1-x_j}\right)^\gamma}{k + \left(\frac{x_j}{1-x_j}\right)^\gamma}.$$

Our goal is to detect a symmetric quintic pitchfork in the scalar equation

$$g(x_1, \mu, k, \gamma, b) = 0.$$

Since the system is Z_2 symmetric with respect to $x_1 = 0.5$ (that is, the change of variables $x_1 \mapsto 1 - x_1$ leaves Eq. (S4) invariant), this singularity has codimension 1 ([2], Theorem VI.5.1(3)). Thus, we only need one unfolding parameter. For analytical tractability and in line with the main text, we fix $b = 0.04$, $\gamma = 2.5$ and use k as the unfolding parameter.

Following the recognition problem ([2], Table VI.5.3), we seek μ^*, k^* such that

$$\begin{aligned}g(0.5, \mu^*, k^*, 2.5, 0.04) &= g_{x_1}(0.5, \mu^*, k^*, 2.5, 0.04) = g_{x_1 x_1}(0.5, \mu^*, k^*, 2.5, 0.04) = 0, \\ g_{x_1 x_1 x_1 x_1}(0.5, \mu^*, k^*, 2.5, 0.04) &\neq 0 \neq g_{x_1 \mu}(0.5, \mu^*, k^*, 2.5, 0.04).\end{aligned}$$

We first solve $g = 0$ in terms of μ and plug the solution into the equation $g_{x_1 x_1 x_1} = 0$ so that it becomes a function only of k . The equation $g_{x_1} = 0$ is automatically solved by picking $x_1 = 0.5$. Solving $g_{x_1 x_1 x_1} = 0$ explicitly in terms of k is non-trivial. Instead, we examine its evolution as a function of k in Fig. S3. The graphical analysis reveals the presence of an isolated zero for $k = k^* \simeq 0.257$. We can now also easily verify that

$$g_{x_1 x_1 x_1 x_1}(1/2, \mu^*, k^*, 2.5, 0.04) < 0, \quad g_{x_1 \mu}(1/2, \mu^*, k^*, 2.5, 0.04) > 0,$$

which completes the recognition problem for the symmetric quintic bifurcation.

S.5 Oscillatory behavior from symmetric quintic bifurcation

Unfolding the symmetric quintic bifurcation illustrated in Fig. 6 in the main text, and proved in Section S.4 in the direction of the subcritical pitchfork, introduces a new kind of multi-stability and with it a hysteresis between the distributed solution and the dominant solutions. This allows for fast switching behavior in the x_i for the closed-loop system with $N = 2$ described by Eqs. (1)-(4) from the main text even as μ increases and decreases slowly through the bifurcation point.

In the case of $N = 2$, $b_{11} = b_{22} = 1$, and $b_{12} = b_{21} = b \in (0, 1)$, Eqs. (1)-(4) specialize to

$$\begin{aligned}\dot{x}_1 &= x_1(\sigma_{\gamma,k}(x_1) + b\sigma_{\gamma,k}(1-x_1))(1-\mu) + (1-x_1)(b\sigma_{\gamma,k}(x_1) + \sigma_{\gamma,k}(1-x_1))\mu - \phi x_1 \\ \dot{\mu} &= -K(x_1 - \alpha_1)(\alpha_2 - x_1)\mu(1-\mu).\end{aligned}\quad (\text{S5})$$

When the two thresholds $x_1 = \alpha_1$ and $x_1 = \alpha_2$ intersect the critical manifold

$$\mathcal{M}_0 = \{(x_1, \mu) : x_1(\sigma_{\gamma,k}(x_1) + b\sigma_{\gamma,k}(1-x_1))(1-\mu) + (1-x_1)(b\sigma_{\gamma,k}(x_1) + \sigma_{\gamma,k}(1-x_1))\mu - \phi x_1 = 0\}$$




Figure S3: Evolution of $g_{x_1 x_1 x_1}$ as a function of k .

along its unstable branches, the closed-loop system (S5) exhibits two distinct oscillatory behaviors (Fig. S4). Each of them corresponds to the periodic alternation between the distributed solution and one of the strongly dominant solutions. The associated limit cycles are symmetric with respect to the line $x_1 = 0.5$. They exist for sufficiently small K .

The proof of existence and stability of these cycles uses techniques from geometric singular perturbations and blow-up theory [3]. In Fig. S5 we provide intuition on how these cycles can be geometrically constructed in the limit $K \rightarrow 0$ by showing the singular ($K \rightarrow 0$) phase portrait. Because x_1 is much faster than μ , trajectories spend most of the time on the critical manifold. The singular limits of the two limit cycles are constructed as closed singular trajectories, which merge along the horizontal part of the critical manifold where $x_1 = 0.5$. Because μ is decreasing in that region, both singular cycles approach the vertical line $\mu = 0$. There, they split in upward and downward directions, respectively. At the intersection with the upper and lower branches of the critical manifold, where μ is increasing, they slide along the critical manifold until the fold singularity, where they jump back to the horizontal branch of the critical manifold, much in the same way as a standard relaxation oscillator.

For small $K > 0$ the two singular cycles perturb to two exponentially stable limit cycles, corresponding to the two oscillatory behaviors in Fig. S4. By Fenichel theory [4], these two cycles are $\mathcal{O}(e^{-1/K})$ -close to each other in the region where they shadow the horizontal branch of the critical manifold. It follows that, for K sufficiently small, tiny perturbations, such as the noise described in the main text, make it possible for the system to switch between the two cycles.

References

- [1] D. Pais, C. H. Caicedo-Núñez, and N. E. Leonard (2012). “Hopf bifurcations and limit cycles in evolutionary network dynamics,” *SIAM Journal on Applied Dynamical Systems*, vol. 11, no. 4, pp. 1754–1784.
- [2] Golubitsky, M. and Schaeffer, D.G. (1985). *Singularities and Groups in Bifurcation Theory*, volume 51 of *Applied Mathematical Sciences*. Springer-Verlag, New York, NY.
- [3] O. Druet, E. Hebey, and F. Robert (2004). *Blow-up Theory for Elliptic PDEs in Riemannian Geometry (MN-45)*. Princeton University Press.
- [4] N. Fenichel (1979). “Geometric singular perturbation theory for ordinary differential equations,” *Journal of Differential Equations*, vol. 31, no. 1, pp. 53–98.




Figure S4: Co-existence of two stable limit cycles (red and blue), shown on the bifurcation diagram for dynamics of Eq. (S5). Stable manifolds are solid lines and unstable manifolds are dashed lines. Parameters are $\gamma = 2.5$, $k = 0.01$, $K = 0.3$, $\alpha_1 = 0.25$, and $\alpha_2 = 0.75$.




Figure S5: Geometric construction of the two stable limit cycles (red and blue) in the limit $K \rightarrow 0$, shown on the bifurcation diagram for dynamics of Eq. (S5). Stable manifolds are solid lines and unstable manifolds are dashed lines. Parameters are $\gamma = 2.5$, $k = 0.01$, $\alpha_1 = 0.25$, and $\alpha_2 = 0.75$.