Collective Performance: Modeling the interaction of habit-based actions

Supplementary Materials

Version 2: 10 September 2010

Massimo Warglien, Universitá Ca Foscari de Venezia, [warglien@unive.it]
Michael D. Cohen, University of Michigan, [mdc@umich.edu]
Daniel A. Levinthal, University of Pennsylvania [dlev@wharton.upenn.edu][*]

[*] The delight of this collaboration has been each author's firm belief that the others made the essential contributions. Accordingly, the order currently shown does not represent an equilibrium. We are grateful for the contributions of many colleagues to the ideas presented here, and especially to Markus Becker, Phanish Puranam, and Jason Owen-Smith for comments on an earlier draft.

Table of Contents

S1. Modeling The Coordination of Perceptual Categories

As a first model instantiating our framework, we introduce a simple representation of the situation space and of how agents categorize it (premises A-E). We show how the model, despite its simplicity, can capture the inherent flexibility of habits and their sensitivity to context. We also show how agents, despite having different individual prototypes of exemplary cases or situations, can learn through minimal changes to tune their perception of the situation space and to achieve greater coordination – a phenomenon clearly underlying the development of successful collective performance (premise L).

Our model could be conceived as a simple example of a model of "appropriate behavior" (March and Olsen 1989), i.e. behavior triggered by the recognition of the type of situation in which an agent finds herself. In this first example, we will concentrate mostly on the first stage of the problem: how a situation can be identified so that the corresponding appropriate action can be performed? (See also Gavetti and Warglien 2008) We will assume that actions are tied to situation categories in a one-to-one correspondence: an assumption that we will relax in later examples.

Thus, the core of this first model is the representation of the situation space and of categorical prototypes within it. We will assume that the "situation space" can be represented literally as a space – actually the product space of multiple original quality dimensions or properties that characterize a set of situations (see section 4.1 in the text). A well-studied psychological example is the space of perceived colors (Gärdenfors 2000); in the economic context one might consider the space representing possible combinations of product fearues.

A single situation can be located as a point in such space. Similarity between situations is defined (following a broad psychological literature) in terms of spatial proximity: the closer the points representing two situations, the more similar to each other they are perceived to be (Shepard 1987) (Nosofsky 1988)[1].

We further assume that agents tend to think coarsely about the situation space, aggregating broad sets of situations that have “family resemblance” in categories (Rosch and Mervis 1975). Categories are often structured around prototypes (Rosch and Mervis 1975), i.e. cases which are “best examples”, or initial examples, of the category. In our model, a prototypical situation is a case highly representative of a class of situations judged by an agent as being “of the same type”. Once they are imprinted in memory, prototypes tend to generate quite automatically (and with reduced memory load) a decomposition of the situation space into categories: a situation tends to be attributed to the category whose prototype is most similar to it.

A simple and elegant geometric model (Gärdenfors 2000) captures the essence of the above statements. If prototypes are points in the situation space and a category is defined as the set of situations to which a given prototype is most similar, the situation space will be naturally decomposed into (convex) regions that represent categories of situations around their respective prototypes (Figure S.1). This decomposition is called a Voronoi diagram, or Voronoi tessellation (Okabe et al. 1992).

---

[1] We do not wish to enter here the controversy on similarity measures originating with Tversky (1977). For a critical appraisal and a defense of geometric measures of similarity, see Gärdenfors (2000).

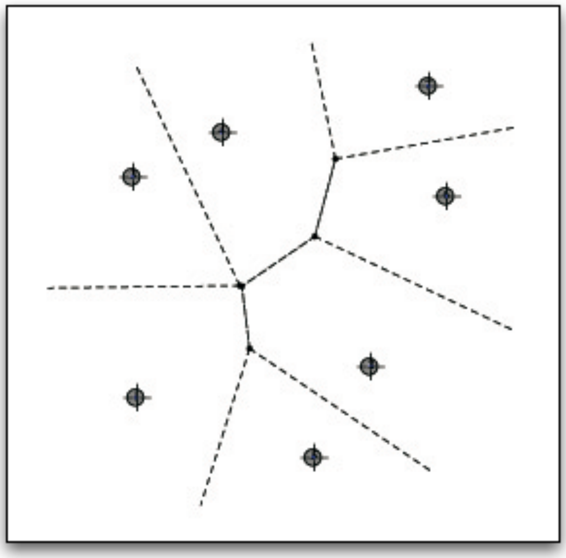

Figure S.1. A Voronoi diagram

Thinking by categories helps to explain the inherent flexibility often observed in habitual behavior and routines: as long as situations are perceived as similar, cases never actually seen previously can be smoothly handled by recognizing the category to which they belong. For example, credit authorization routines for customers with similar measures of credit worthiness will trigger similar actions that in turn generate similar loan contracts.

Moreover, it is possible to introduce much more subtle forms of flexibility by minimal extensions of this basic model. It has been repeatedly observed (Barsalou 1987) that when engaged in similarity judgments, individuals do not give unvarying weights to the different attributes along which similarity is assessed. Instead, they tend to make an adaptive use of attribute weighting. For example, in tasks in which one must discriminate between two categories, individuals tend to emphasize or increase the salience of attributes along which prototypical exemplars are most different (Nosofsky 1988).

Our geometrical model of the situation space can easily capture this adaptive use of salience by introducing weighted dimensions of similarity (Nosofsky 1988) (Gärdenfors 2000). If the weight of the dimensions of the situation space can be changed, distance and thus similarity between situations will change accordingly, and the whole decomposition of the situation space into categories will be modified.[2]

Figure S.2 provides a simple two-dimensional illustration of this effect. Five prototypical points define categories that shift in space as the relative weight (w) of the left right dimension increases. As can be clearly seen, the attribution of a situation (the black dot) to categories shifts accordingly.

The sensitivity of the categorization of situations to salience changes suggests a very simple process that makes similarity judgment – and thus action triggering – responsive to the context, providing further smoothness to habitual behavior and illustrating the "tunable" nature of the action function (see section 4.2 in the main text). Indeed, salience tends to respond to important environmental cues. For example, seeing a car accident is likely to emphasize an individual's perception of the risk dimension of driving situations, modifying the classes of situations in which a driver finds it appropriate to reduce speed. In general, sensitivity to context via salience will produce a typical fuzziness of categorical attribution in the "peripheral" regions of categories (situations which are located close to the boundaries), while the category membership of situations close to prototypes will tend to remain stable. It is important to point out that these changes can happen without any modification of the categorical prototypes

[2] Those familiar with the principles of geometry will recognize this as simple affine transformation of the state space.

– just as a result of shifts in the salience of dimensions. They may also be largely unconscious, as when changes in salience are triggered by emotions.

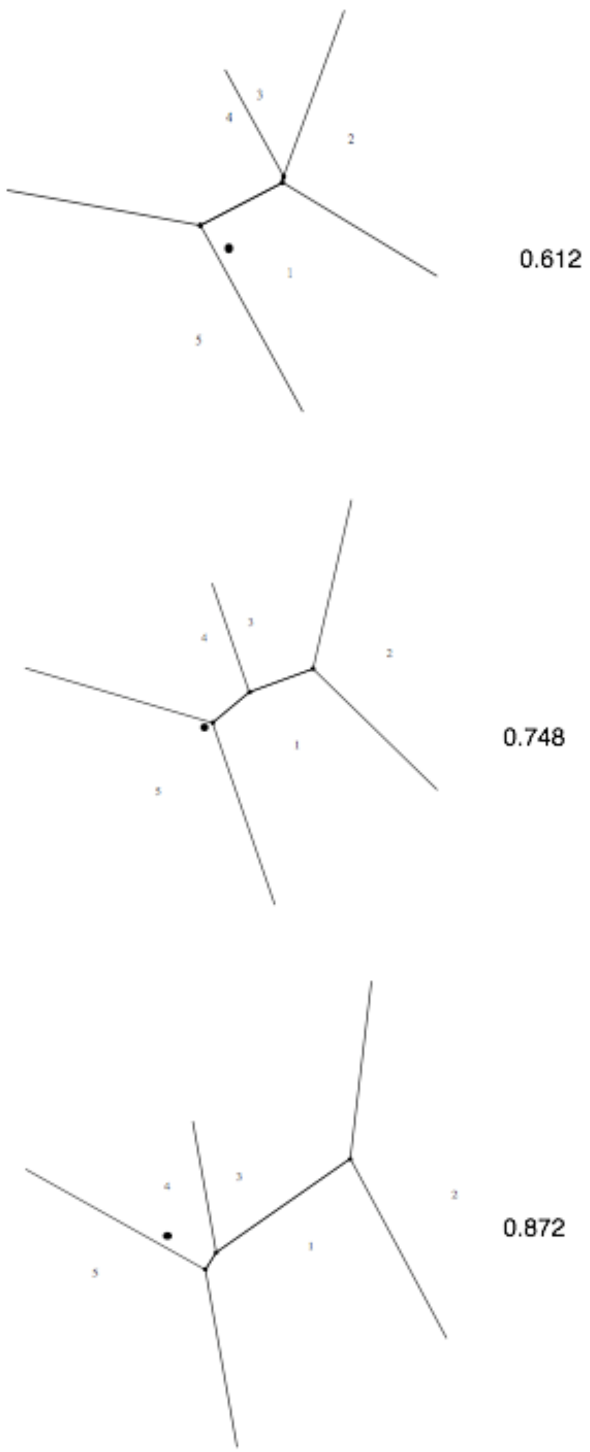

Figure S.2: Voronoi diagrams of space with five prototypes (shown by numbers). The situation (black dot) is categorized differently at each of three different levels of relative weight w for the left-right dimension.

The model above can be extended to multiple agents. Indeed, it allows us to capture two sources of heterogeneity in their perception of situations: differences in prototypes, and differences in the weight accorded to attributes. Some weights might be equal to zero, bringing within the scope of the model instances where agents consider different dimensions of a situation.

Clearly, differences in the way situations are categorized may be a fundamental source of coordination failures. We show that a basic form of incremental salience adjustment may markedly improve coordination between agents who hold heterogeneous perceptions of the situation space in the absence of any modification of their prototypes.

For sake of simplicity, imagine a problem in which there are two agents, and two actions (respectively: {A, B} and {a, b}) available to each one of them. Actions are complementary, so that they produce a coordinated outcome only when their combination is <A, a>, both agents choose the first of their two actions, or <B, b>, both agents choose the second of their two actions. Furthermore, agents have two prototypes each, which are related to actions. So when the situation is perceived as being of type α, agents will trigger, respectively, actions A and a; alternatively, when the situation is perceived as being of type β, agents will trigger, respectively, actions B and b. The problem arises from the fact that the prototypes, and therefore the category boundaries, for the situations they label 'α' and 'β,' may be different between the two agents (see Figure 3). Furthermore, agents may start with different weights for the attribute dimensions of the situation space.

Now, imagine that situations are encountered by the two agents, and at each new situation they act independently and observe whether their actions were coordinated or not. Whenever coordination failures are experienced, each agent will react by dampening the relative weight (salience) of the dimension that contributed most prominently to the coordination failure.

Figure S.3 shows how coordination failures due to the initial lack of alignment of individual categories are corrected by this simple process of "salience tuning". The graph shows the evolution of the average number of coordination failures as 500 different random situations are sequentially shown to a pair of agents, with substantial improvement occurring quite quickly.[3]

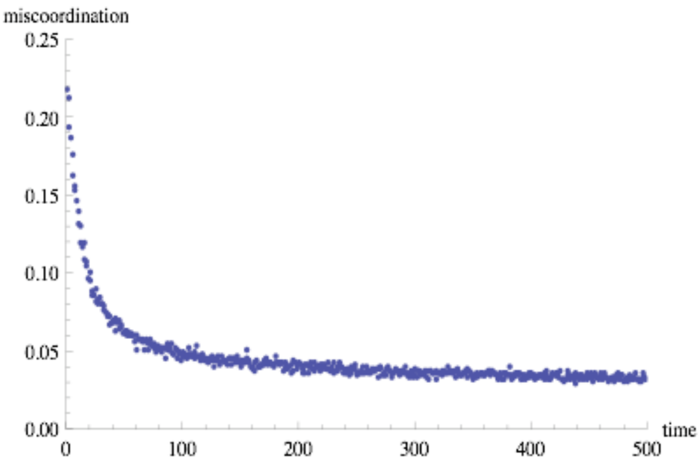

Figure S.3 Decreasing coordination failures with salience tuning.

Figure S.4 compares the initial and final categorizations of the space in a single sequence of a pair that has learned over 500 different situations. The alignment of categories (despite the persistent heterogeneity of the prototypes)

[3] Averages are computed over 100 different random sets of pairs of prototypes and initial weights of the space dimensions. In turn, for each set of prototypes and weights, 100 different series of 500 random situations are run.

can be clearly seen.[4] Further notice the rather rapid mutual adaptation process: after 20 trials coordination succeeds in about 90% of the cases (10% of errors), and after 80 trials the success rate reaches approximately 95%.

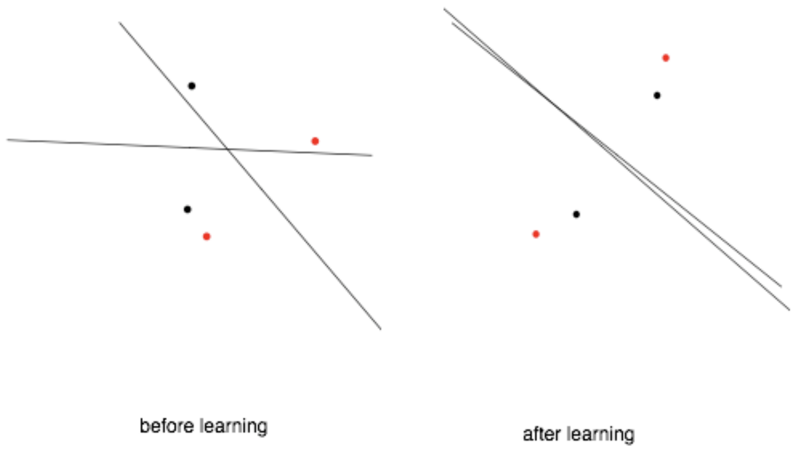

Figure S.4: Co-adaptation of two actors' conceptual spaces.

[4] Like most learning processes, though, this one will work only within some boundary conditions. It requires that the sets of prototypes of the agents be ordered in the same way on any dimension (order-monotonicity). For example, prototypes of a category X should be North-East of those of a category Y for both agents. In other words, the model suggests the process will work, even if the prototypes are in different locations for the two actors, so long as they are similarly ordered on the dimensions that define the category space. This seems a reasonable requirement that individuals share some common structure in their representations that in turn may subsume a common coarse structure of experience.

Supplementary Model S2:

Modeling acts and actions in 'transform-the-target'

As a second illustration of modeling within the framework we consider a two-player experimental task, the 'Transform-The-Target' card game which has been used to establish key features of routinized activity in a controlled laboratory setting. Multiple experiments have shown that dyadic, habit-based action patterns develop within the first few deals, after just a couple minutes of play (Cohen and Bacdayan 1994; Egidi and Narduzzo 1997; Wang and Zhang 2008).

This experiment has proven to be a useful context in which to study the emergence of recurring interaction patterns among individuals. The emergent behaviors have many of the properties associated with skillful collective performance, including the rapid execution of interdependent behavior and "suboptimality", the property that while patterns emerge to achieve the desired goal, they do not necessarily provide the most efficient solution path to the problem the actors face.

Background Task Information

We begin with a presentation of the task details for those not familiar with Transform-The-Target (TTT). For the two-person version, the game is played with six cards, the 2, 3, and 4 of hearts and of spades[1]. The board, as shown in Figure S2.1, has

---

[1] Versions for more than two players have also been developed (Wang and Zhang 2008; Wollersheim 2009).

four positions while each of the two players holds one card in the positions labeled HAND-CK, at the top, and HAND-NK, at the bottom.[2] The key board position is the TARGET. The aim of play is to maneuver the red 2 into the TARGET area. Among the other three board positions, two are occupied by cards that are placed face-down (the DOWN-A and DOWN-B positions), and the third holds a card that is face-up (the UP position), and therefore visible to both players. Play proceeds in alternating turns, beginning with the player labeled 'ColorKeeper'. A turn allows a player to exchange his or her HAND card with one of the board cards, or to "pass". When a series of exchanges with the board successfully moves the red 2 into the TARGET area, the play for that deal of the cards ends. No verbal communication is allowed. In a typical experiment, a dollar might be won for completing the deal. The number of moves by both players would be counted, and each move might cost $.10. After 40 deals of the cards are played, each deal providing a different starting configuration, the net earnings would be divided between the two players.

What makes the game challenging is one further rule. The players do not have symmetric capabilities. One player, the one designated "ColorKeeper", can exchange with the TARGET only when that player's hand card, HAND-CK and the TARGET card are of the same color. The other player, designated "NumberKeeper", can exchange with the TARGET only when the HAND-NK and TARGET cards are of the same number.

---

[2] Board positions are labeled with capital letters.

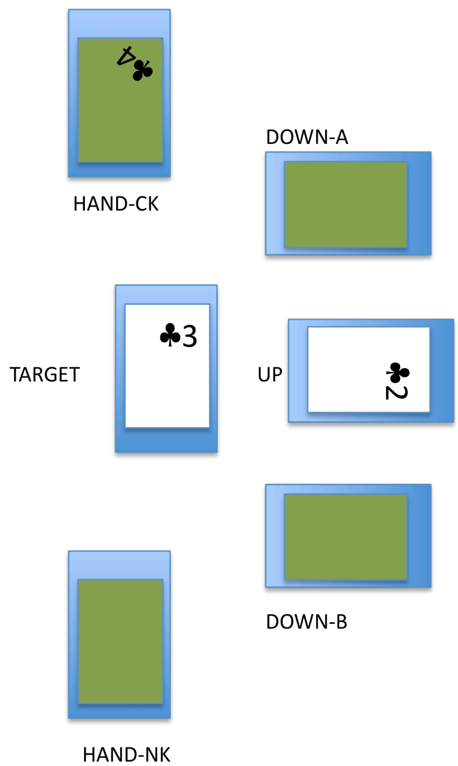

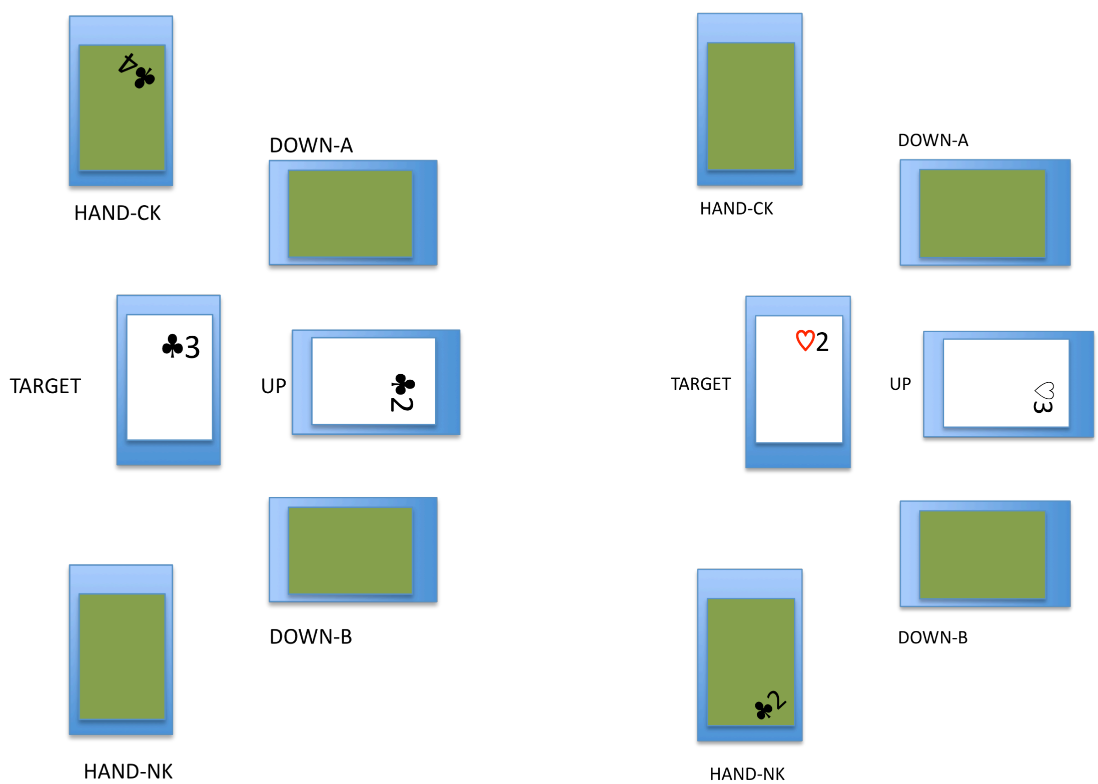

Figure S2.1. Left panel (a) showing cards visible to ColorKeeper at start of deal. Right panel (b) showing configuration of cards after NumberKeeper's finishing move.

In TTT there are 720 permutations of the six playing cards, and there are two players who might have the next turn, so the game can be in 1440 states. Of these, 240 correspond to the goal, that is, the red 2 is in the TARGET. This state space is small compared to real world problems, but big enough to keep novice players quite challenged for the duration of a 40-minute experimental session. The successive deals the players encounter, which come in a predetermined sequence, present a number of difficulties. Some are quite easy, and can be solved in just 2 or 3 smart moves. Typical hands require 4 or 5 moves. Some are quite difficult, and even good players may need 6 or 7 moves. It is not unusual for beginners to slowly take 10 or 15 steps to complete a deal that players with well-developed action patterns can do quickly in 4 or 5.

Row Nmbr	To Move	HndCK	DownA	UP	DownB	HndNK	TARGET	Move Made
1	CK	4♣	3♥	2♣	2♥	4♥	3♣	DownB
2	NK	2♥	3♥	2♣	4♣	4♥	3♣	DownA
3	CK	2♥	4♥	2♣	4♣	3♥	3♣	Up
4	NK	2♣	4♥	2♥	4♣	3♥	3♣	Up
5	CK	2♣	4♥	3♥	4♣	2♥	3♣	Target
6	NK	3♣	4♥	3♥	4♣	2♥	2♣	Target
7	Done	3♣	4♥	3♥	4♣	2♣	2♥	

Table S2.1: Successive states of play in example hand of TTT. Columns indicate the player to move next, the cards occupying each of the six positions, and the move made by the player to produce the state in the following row. Underlines indicate the cards that are exchanged in producing the state in the following row.

It may help to build intuition for the patterned action we are modeling if we "walk through" play of a typical deal as listed in Table S2.1. Figure S2.1 and the first line of the Table show the positions of cards just after our example hand has been dealt. The player in the "ColorKeeper" (CK) role has the first move. She sees that the red 2 is neither in her hand, HAND-CK, which holds the black 4, nor in the TARGET or UP positions of the board. (The two DOWN cards, -A and -B, and the NumberKeeper's HAND-NK card cannot be seen by CK.) CK searches for the red 2 by exchanging her HAND-CK card with the card at board position DOWN-B. In our example CK is fortunate in her search and now holds the red 2, and it is NK's turn. We are at row 2 of the Table. He exchanges his HAND-NK with DOWN-A (underlined). He is looking for the red 2 as well, and, since he can't see DOWN cards or HAND-CK, he doesn't know CK has found the red 2.

CK cannot finish at row 3 by putting the red 2 into the TARGET since that would change the card color in the TARGET, violating the definition of the ColorKeeper role. (For that matter, NK couldn't put the red 2 in if he held it in HAND-NK either. He needs the TARGET to contain the black 2 in order for that exchange to keep the TARGET number unchanged, as required by his role.) CK's next act is exchanging her hand card with the UP card on the board. (Yielding row 4.) After that, experienced players would typically be quick to complete the hand. NK also exchanges with UP. CK puts the black 2 in the TARGET. NK puts in the red 2, and the hand is finished after six moves. (Table S2.1, row 7, and Figure S2.1b.)

This is not actually the ideal solution for this particular deal of the cards. For example, if CK exchanges with UP on the first move, there is a good chance of finishing in four steps instead of six. But pairs of players develop their action patterns over a number of different deals they face in their early experience, and then apply them to get good results in a new situation, sometimes without noticing that still better options existed. Such occasions of smooth-though-sometimes-suboptimal action are a hallmark of habit-based action patterns. The success of TTT as an experimental instrument has been its ability to evoke just such patterns from subjects.

Presentation of the Model

We show that play of the game can be represented as a computational model embodying many of the premises of section 2.1 of our main text. The ultimate goal of the game is to place a specific card (the 2 of hearts) in a specific location on the board (the TARGET) subject to constraints of information (hidden cards) and action (specialized player roles). The constraints on action also preclude any explicit communication among the players. While placing the 2 of hearts in the TARGET location is the ultimate goal, players generally develop sub-goals or intermediate ends-in-view, such as placing a card in a jointly visible area that enables action on the part of the other player. The most basic actions (premise I) of the TTT game are each player's exchanges of the card in their hand with four possible board positions, or "passing" their turn.[3] Each player therefore has

---

[3] This model does not incorporate the manner of performing each action, although human players do move physical playing cards, or drag card-images on computer screens.

five fundamental *actions*. The game aligns with the most fundamental aspect of the paper's framework: each action is a *function* that transforms the game-play *situation* and is applicable only in certain conditions (C). For an action to be available to a player, it has to be the player's turn, and, in the case of exchanges with the TARGET, the cards in that position and in the player's HAND have to jointly conform to the player's role restriction. Each action (except pass), changes three aspects of the TTT world: a board position gets a new card; the player's HAND position gets the card from the board position; and it becomes the other player's turn to move. Each action is potentially applicable in a large set of board situations, since the remaining cards can be in many different permutations among the unaffected board positions.

Each action might be invoked in the pursuit of several different near-term objectives, or *activated-ends* (F). For example, exchanging with one of the DOWN cards might be done in an effort to locate the red 2.[4] But at other times, a player might exchange for a DOWN card as part of seeking a card needed to prepare the TARGET for a subsequent finishing move.

The program fragment in Figure S2.2 shows how these components of an action can be incorporated into a simulation model.[5] The function definition is shown with 3 inputs: the current state of game play, *situation_* ; a predicate (returning True or False) that defines a desired condition for the future, *endInViewQ_* ; and a possible board transformation, *act_* . The function operates on the board if the state of the board is

---

[4] The game is played with six cards: the 2, 3, 4 of hearts and spades. Therefore, it is sufficient to distinguish the suits by their color and the computational model refers to the cards as r2, r3, r4, b2, b3, b4.

[5] The language used is Mathematica (v6). Many other languages would be possible, of course. Mathematica has been used in view of its strength in combining functional and declarative (rule-based) programming styles. It is fundamental to our framework to represent action as a function and Mathematica provides a medium in which this approach is easily expressed.

consistent with the action's conditions, and the expected result of acting (E, L) would achieve the condition desired, the activated-end. For brevity, the ancillary functions *expectations*, *transform*, and *adapt* are not shown.

```
…

In[4]:=
action[situation_, activatedEndQ_, act_] :=
    Module[{s = situation, eQ = activatedEndQ, a = act},
        If[eQ[expectation[s, a]],
            {transform[s, a], adapted[eQ, s, a, True]},
            {s, adapted[eQ, s, a, False]}
    ]
];

In[5]:=
{resultSituation, resultActivatedEnd} =
    action[{r2, r3, b2, b4, r4, b3, ck}, ckHandB2Q, exchangeUp]

Out[1]=
  {{"b2", "r3", "r2", "b4", "r4", "b3", "nk"}, {ckHandB2Q, "attained"}}
```

Figure S2.2: Mathematica code for a function implementing *action* for the game TTT.

Also shown in Figure S2.2 is an input line, In[5], invoking the function in a particular *situation_* that is shown as a sequence of cards and an indication of which player has the next move that occupy the ordered positions {HAND-CK, DOWN-A, UP, DOWN-B, HAND-NK, TARGET, next-mover}. The *activatedEndQ_* invoked in this case is ckHandB2Q, a predicate that tests "Is the HAND-CK card the black 2?" The *act_* under consideration is exchangeUp, an exchange of HAND-CK with the UP board position. The action will return two results. The transformation is activated in this case since the *expectations* associated with the situation and action are consistent with the activated-end. Therefore, the function returns a list of two items, shown as Out[1]: (1) the *resultSituation,* which reflects the *transform* of the board positions by exchanging the cards in the UP and HAND-CK positions and advancing the next-mover, and (2) a list indicating the *activatedEnd* has been attained as the result of the function *adapted* which has been called with a parameter of True that indicates an act occurred. Thus, the action function has produced an act (G). If the transformation were not activated, the action function would have produced no act, and would have returned two different results: (1) the unchanged board situation with unchanged next-mover, and (2) a list indicating a possible revision of the *activatedEnd* resulting from the function *adapted* called with a parameter of False.

The action uses the function *adapted* to check whether its results have met its expectations, whether it has produced the desired result, or whether an inability to act signals a need for changed objectives (F, K). It uses these indicators of its performance to modify activated-ends using an updating process that is not shown here in detail. This updating allows the activated-ends to be strengthened or generalized if the action achieves expectations and activated-ends, and to be repaired if there has been a *breakdown*.(B,K). [6]

Our discussion of this computational model of TTT has not exercised the model in full play conditions. The complexity of that exposition is beyond the scope of this paper. Instead, we have described only a key portion of the model in order to demonstrate

---

[6] In the conditions of this example, the first order expectations of these lowest level actions will always be met. The program is not allowed ever to mistakenly alter the board into an impossible configuration or lose track of which player will move next. For higher order actions, however, it is possible that expectations will not be met, mainly because of imperfect generalizations from prior experience.

how natural it is to embody the premises of our framework in a small computer program and to represent recurring action patterns as functions.

Works Cited

Cohen, M. D., and P. Bacdayan

1994 "Organizational Routines Are Stored as Procedural Memory: Evidence From a Laboratory Study." Organization Science 5:554-568.

Egidi, Massimo, and Alessandro Narduzzo

1997 "The emergence of path-dependent behaviors in cooperative contexts." International Journal of Industrial Organization 15:677-709.

Wang, Jian-An, and Gang Zhang

2008 "Knowledge, Routines and Performance in Collective Problem Solving." Acta Pssychologia Sinica 40:862-872.

Wollersheim, Jutta

2009 "An Empirical Investigation of the Role of Dynamic Capabilities in Mergers and Acquisitions." Rotterdam, Netherlands.