

The Logic of Doxastic Strategies

Junli Jiang,¹ Pavel Naumov²

¹ Institute of Logic and Intelligence, Southwest University, China

² University of Southampton, United Kingdom
walk08@swu.edu.cn, p.naumov@soton.ac.uk

Abstract

In many real-world situations, there is often not enough information to know that a certain strategy will succeed in achieving the goal, but there is a good reason to believe that it will. The paper introduces the term “doxastic” for such strategies. The main technical contribution is a sound and complete logical system that describes the interplay between doxastic strategy and belief modalities.

1 Introduction

During the 1991 Gulf War, American troops successfully used surface-to-air Patriot missiles to intercept Iraqi-modified Scud tactical ballistic missiles originally designed by the Soviet Union in the 1950s. In fact, on 15 February 1991, President George H. W. Bush travelled to the Patriot manufacturing plant in Andover, Massachusetts to praise the missiles’ capability:

You see, what has taken place here is a triumph of American technology . . . we are witnessing a revolution in modern warfare, a revolution that will shape the way that we defend ourselves for decades to come. For years, we’ve heard that antimissile defenses won’t work, that shooting down a ballistic missile is impossible – like trying to “hit a bullet with a bullet.” Some people called it impossible; you called it your job. They were wrong, and you were right . . . Patriot is 41 for 42: 42 Scuds engaged, 41 intercepted. (Bush 1991)

Just 10 days after this speech, an American radar detected an Iraqi Scud approaching Dhahran, Saudi Arabia.

Traditionally, knowledge is modelled through an indistinguishability relation that captures an agent’s own observations. However, modern autonomous agents often obtain their knowledge from the data that they receive from other systems. This was the case on 25 February 1991. By analysing the position, \vec{x} , and the velocity, \vec{v} , of the Scud, the Patriot concluded that intercepting the Scud was within its capabilities. We write this as

$$K_{\vec{x},\vec{v}}(\text{“The Patriot has a strategy to destroy the Scud”}) \quad (1)$$

and say that dataset $\{\vec{x}, \vec{v}\}$ informs the knowledge of the statement “The Patriot has a strategy to destroy the Scud”. In general, we write $K_X\varphi$ if knowing the values of all variables in set X is enough to conclude that statement φ is true. We refer to sets of data variables, like X , as *datasets*. Modality K_X for a set X of *Boolean* variables is introduced by Grossi, Lorini, and Schwarzentruher (2015). For arbitrary datasets, it is proposed by Baltag and van Benthem (2021). The name “data-informed knowledge” is suggested by us in (Jiang and Naumov 2022a). We also proposed a dynamic logic for this modality (Deuser et al. 2024).

Of course, for the Patriot to know that it has a strategy is not the same as to know what the strategy is. The distinction between “knowing that a strategy exists” and “knowing what the strategy is”, in the case of the traditional (indistinguishability-relation-based) knowledge, has been studied in various logics of know-how strategies (Ågotnes and Alechina 2019; Naumov and Tao 2017; Fervari et al. 2017; Naumov and Tao 2018c,b,a; Fervari, Velázquez-Quesada, and Wang 2021; Li and Wang 2021).

In our data-informed setting, in order for the Patriot to know the strategy (direction in which to launch a surface-to-air missile), it also must know the exact time, t , since the position, \vec{x} , and the velocity, \vec{v} , of the Scud has been measured. We write this as

$$[\text{Patriot}]_{\vec{x},\vec{v},t}(\text{“Scud is destroyed”}) \quad (2)$$

and say that the dataset $\{\vec{x}, \vec{v}, t\}$ informs the strategy of Patriot to achieve the condition “Scud is destroyed”. In general, we write $[C]_X\varphi$ if knowing the values of all variables in dataset X is enough to know a strategy for coalition C to achieve condition φ . We introduced this modality in (Jiang and Naumov 2022a).

The values in the dataset $\{\vec{x}, \vec{v}, t\}$, obtained from the radar, gave the Patriot enough information to compute the strategy (direction of missile launch) that was guaranteed to destroy the Scud. Upon computing the strategy, the Patriot sent it to the launcher. . . Seconds later, Scud hit the barracks of the 14th Quartermaster Detachment of the US Army’s 99th Infantry Division.

To understand what went wrong that day in the skies above Dhahran, Saudi Arabia, one needs to acknowledge the fact that data does not always correctly reflect the state-of-affairs of the real world. Even if the Patriot knows the “true”

values of \vec{x} and \vec{v} as measured by the radar, these values might not reflect the correct position and velocity of Scud due to, for example, failure in the radar operation. In the case of our statement (1), it is perhaps more sensible to say that by analysing the radar-provided values of \vec{x} and \vec{v} and *trusting* those values, the Patriot formed a *belief* that intercepting the Scud is within its capabilities:

$$B_{\vec{x}, \vec{v}}^{\vec{x}, \vec{v}}(\text{“The Patriot has a strategy to destroy the Scud”}). \quad (3)$$

In general, we write $B_X^T \varphi$ if under the assumption of the trust in dataset T , the values of variables in the dataset X inform belief φ . We proposed this *trust-based belief* modality in (Jiang and Naumov 2022b).

As it turns out, the statement $B_X^{\mathcal{Q}} \varphi$ is equivalent to the statement $K_X \varphi$. In other words, any belief, which is not based on trust, is knowledge. Any such belief, of course, is true. In general, just like the other types of beliefs, trust-based beliefs are not always true. They might be false if they are based on *trust* in *non-trustworthy* data. This, however, was *not* the case on 25 February 1991. The position and the velocity of the Scud communicated by the radar correctly reflected the physical characteristics of the approaching ballistic missile. Thus, the Patriot indeed had the capability to intercept the Scud.

As we have just observed when the trust in data is taken into account, the knowledge statement (1) becomes a belief statement (3). Of course, a similar adjustment should be made to statement (2). The Patriot did not actually *know* the strategy to destroy the Scud, it only had a *belief that a particular strategy will work*. This belief was based on the *trust* in the dataset $\{\vec{x}, \vec{v}, t\}$. We write this as:

$$[\text{Patriot}]_{\vec{x}, \vec{v}, t}^{\vec{x}, \vec{v}, t}(\text{“Scud is destroyed”}). \quad (4)$$

In general, we write $[C]_X^T \varphi$ if, under the assumption of trust in dataset T , dataset X informs a belief about some strategy of coalition C that it will achieve condition φ . We call such strategies “doxastic” (related to beliefs). Doxastic strategies are guaranteed to work if dataset T is not only *trusted* but also *trustworthy*.

The Patriot system clock stored time (since boot) measured in tenths of a second. The actual time was computed by multiplying the stored value by 0.1 using a 24-bit fixed point register. This introduced an error of about 0.00000095 seconds. By the time the Scud appeared in Dhahran, the system had been running for over 100 hours, creating an accumulated time difference between the radar’s and the launcher’s clocks of about 0.34 seconds. A Scud travels over 500 meters in this time (Mignotte 2010). The strategy that the Patriot *believed* would destroy the Scud was based on *trusted* but *non-trustworthy* time data. The “triumph of American technology” pointed the missile launcher in a completely wrong direction, killing 28 American soldiers (Mignotte 2010) and wounding close to 100 (Apple 1991). A year later, on 11 November 1992, in Greensburg, Pennsylvania, a memorial was dedicated to the soldiers killed by the Scud attack. The monument is facing in the direction of Dhahran, Saudi Arabia (Legion 2017).

In this paper, we introduce formal semantics for doxastic strategies and a sound and complete logical system that describes the interplay between modalities B_X^T and $[C]_X^T$.

The paper is structured as follows. In the next section, we define the class of games that we use to model multiparty interactions. Next, we define the syntax and the formal semantics of our logical system. In Section 4, we list its axioms and inference rules. Section 5 contains the proof of the completeness. Section 6 discusses *ex ante* and *ex post* trust. Section 7 concludes. The proof of soundness and one of the lemmas from Section 5 can be found in the full version of this paper (Jiang and Naumov 2023).

2 Games

In this section, we introduce the class of games that is used later to give a formal semantics of our logical system. Throughout the paper, we assume a fixed set of atomic propositions and a fixed set of data variables V . In addition, we assume a fixed set of *actors* \mathcal{A} . We use the term “actor” instead of the more traditional “agent” to emphasise the fact that in our setting knowledge and beliefs come from data and they are not related to actors, who are only endowed with an ability to act.

By a *dataset* we mean any subset of V . Informally, given a game, we think about each data variable as having a value in each state of the game. However, formally, it is only important to know, for each two given states, if a variable has the same or different values in those states. Thus, we only need to associate an indistinguishability relation \sim_x with each data variable $x \in V$. Informally, $w \sim_x u$ if data variable x has the same value in states w and u .

To model the trustworthiness of data variables, for each state w of a game we specify a dataset $\mathcal{T}_w \subseteq V$ consisting of the variables that are trustworthy (but not necessarily trusted) in the state w . We introduced this way to model the trustworthiness of data in (Jiang and Naumov 2022b).

Finally, in the definition below and throughout the paper, by Y^X we mean the set of all functions from set X to set Y .

Definition 1 A tuple $(W, \{\sim_x\}_{x \in V}, \{\mathcal{T}_w\}_{w \in W}, \Delta, M, \pi)$ is called a *game* if

1. W is a (possibly empty) set of states,
2. \sim_x is an indistinguishability equivalence relation on set W for each data variable $x \in V$,
3. \mathcal{T}_w the set of data variables that are “trustworthy” in state $w \in W$,
4. Δ is a nonempty set of “actions”,
5. $M \subseteq W \times \Delta^{\mathcal{A}} \times W$ is a “mechanism” of the game,
6. $\pi(p) \subseteq W$ for each atomic proposition $p \in P$.

By a *complete action profile* we mean an arbitrary element of the set $\Delta^{\mathcal{A}}$. By a *coalition* we mean an arbitrary subset $C \subseteq \mathcal{A}$ of actors. By an *action profile* of a coalition C we mean an arbitrary element of the set Δ^C .

The mechanism of the game M specifies possible transitions of the game from one state to another. If $(w, \delta, u) \in M$, then under complete action profile δ from state w the game can transition to state u . Note that we do not require the mechanism to be deterministic. We also do *not* require that

for each state $w \in W$ and each complete action profile δ there is at least one state u such that $(w, \delta, u) \in M$. If such state u does not exist, then we interpret this as a termination of the game upon the execution of profile δ in state w .

3 Syntax and Semantics

Language Φ of our logical system is defined by the grammar

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \rightarrow \varphi) \mid \mathbf{B}_X^T\varphi \mid [C]_X^T\varphi,$$

where p is an atomic proposition, $C \subseteq \mathcal{A}$ is a coalition, and $T, X \subseteq V$ are datasets. We read expression $\mathbf{B}_X^T\varphi$ as “under the assumption of trust in dataset T , dataset X informs belief φ ”. And we read expression $[C]_X^T\varphi$ as “under the assumption of trust in dataset T , dataset X informs a doxastic strategy of coalition C to achieve φ ”.

We write $f =_C g$ if $f(x) = g(x)$ for each element $x \in C$ and $w \sim_X u$ if $w \sim_x u$ for each data variable $x \in X$.

Definition 2 For any formula $\varphi \in \Phi$, and any state $w \in W$ of a game $(W, \{\sim_x\}_{x \in V}, \{\mathcal{T}_w\}_{w \in W}, \Delta, M, \pi)$, the satisfaction relation $w \Vdash \varphi$ is defined as follows

1. $w \Vdash p$, if $w \in \pi(p)$,
2. $w \Vdash \neg\varphi$, if $w \not\Vdash \varphi$,
3. $w \Vdash \varphi \rightarrow \psi$, if $w \not\Vdash \varphi$ or $w \Vdash \psi$,
4. $w \Vdash \mathbf{B}_X^T\varphi$, if $u \Vdash \varphi$ for each state $u \in W$ such that $w \sim_X u$ and $T \subseteq \mathcal{T}_w$,
5. $w \Vdash [C]_X^T\varphi$, when there is an action profile $s \in \Delta^C$ of coalition C such that for all states $u, v \in W$ and each complete action profile $\delta \in \Delta^A$, if $w \sim_X u$, $T \subseteq \mathcal{T}_w$, $s =_C \delta$, $(u, \delta, v) \in M$, and $T \subseteq \mathcal{T}_v$, then $v \Vdash \varphi$.

Note that, by item 4 of the above definition, $w \Vdash \mathbf{B}_X^T\varphi$ states that condition φ is satisfied in each state $u \in W$ indistinguishable from state w by dataset X . Thus, \mathbf{B}_X^T is equivalent to data-informed knowledge modality \mathbf{K}_X discussed in the introduction.

Item 5 above captures our informal intuition that $[C]_X^T\varphi$ means that dataset X informs the knowledge of a strategy that guarantees the achievement of φ in states where dataset T is trustworthy. Note that there are multiple ways to formalise this: we can require T to be trustworthy in state u (before the transition), in state v (after the transition), or in both of these states. By including conditions $T \subseteq \mathcal{T}_u$ and $T \subseteq \mathcal{T}_v$ in item 5, we require T to be trustworthy *ex ante* (before transition) and *ex post* (after transition). We discuss an alternative approach in Section 6.




Figure 1: A game. Data variables \vec{x} and \vec{v} are trustworthy in all states. Data variable t is trustworthy in states w_2, w_3, w_4 .

Figure 1 depicts a (very simplistic) game capturing our introductory game. This game has a single actor, the Patriot.

The actual state is w_1 , in which data variable t is not trustworthy. The only state which is $\{\vec{x}, \vec{v}, t\}$ -indistinguishable from the current state and in which all variables in dataset $\{\vec{x}, \vec{v}, t\}$ are trustworthy is state w_2 . Note that, in state w_2 , strategy s_1 can be used to destroy the Scud. Thus, in the current state w_1 , dataset $\{\vec{x}, \vec{v}, t\}$ informs the belief that the Patriot can use strategy s_1 to destroy the Scud. Hence, in this game, formula (4) is satisfied in state w_1 .

4 Axioms

In addition to propositional tautologies in language Φ , our logical system contains the following axioms.

1. Truth: $\mathbf{B}_X^\emptyset\varphi \rightarrow \varphi$,
2. Negative Introspection: $\neg\mathbf{B}_X^T\varphi \rightarrow \mathbf{B}_X^\emptyset\neg\mathbf{B}_X^T\varphi$,
3. Distributivity: $\mathbf{B}_X^T(\varphi \rightarrow \psi) \rightarrow (\mathbf{B}_X^T\varphi \rightarrow \mathbf{B}_X^T\psi)$,
4. Trust: $\mathbf{B}_X^T(\mathbf{B}_Y^T\varphi \rightarrow \varphi)$,
5. Monotonicity: $\mathbf{B}_X^T\varphi \rightarrow \mathbf{B}_{X'}^{T'}\varphi$ and $[C]_X^T\varphi \rightarrow [C']_{X'}^{T'}\varphi$, where $T \subseteq T'$, $X \subseteq X'$, and $C \subseteq C'$,
6. Cooperation: $[C]_X^T(\varphi \rightarrow \psi) \rightarrow ([D]_X^T\varphi \rightarrow [C \cup D]_X^T\psi)$, where $C \cap D = \emptyset$,
7. Strategic Introspection: $[C]_X^T\varphi \leftrightarrow \mathbf{B}_X^T[C]_X^T\varphi$,
8. Belief in Unavoidability: $\mathbf{B}_X^T[\emptyset]_Y^T\varphi \rightarrow [\emptyset]_X^T\varphi$,
9. Public Belief: $\mathbf{B}_\emptyset^T\varphi \rightarrow [\emptyset]_\emptyset^T\varphi$.

The Truth axiom, the Negative Introspection axiom, the Distributivity axiom, the Trust axiom and the belief part of the Monotonicity axiom are the axioms for the trust-based beliefs as stated in (Jiang and Naumov 2022b). The most non-trivial among them is the Trust axiom. To understand the meaning of this axiom, note that although the principle $\mathbf{B}_X^T\varphi \rightarrow \varphi$ is not universally true, it is true in all states where dataset T is trustworthy. This observation is captured in the Trust axiom.

The Cooperation axiom is a variation of the axiom from the Logic of Coalition Power (Pauly 2002). The Strategic Introspection axiom states that dataset X informs a doxastic strategy iff X informs a belief in having a doxastic strategy. The Strategic Introspection axiom and the Belief in Unavoidability axiom (which we discuss below) can both be stated in alternative forms where the superscript of modality \mathbf{B} is the empty set \emptyset .

To understand the meaning of the Belief in Unavoidability axiom, note that statement $[\emptyset]_X^\emptyset\varphi$ means that dataset X informs the knowledge that condition φ is *unavoidably true* in the next state. Similarly, $[\emptyset]_X^T\varphi$ means that, under the assumption of trust in dataset T , dataset X informs the *belief* that condition φ is unavoidably true in the next state. The Belief in Unavoidability axiom states that *if X informs the belief that Y informs the belief in unavoidability of φ , then X itself informs this belief in unavoidability of φ .*

Finally, observe that $\mathbf{B}_\emptyset^\emptyset\varphi$ means that statement φ is true in all states of the game. In other words, φ is public knowledge in the game. The Public Belief axiom states that if condition φ holds in all words where T is trustworthy, then it is believed to be unavoidably true in the next state.

We write $\vdash \varphi$ and say that formula $\varphi \in \Phi$ is a theorem of our system if it is derivable from the above axioms using the Modus Ponens and the Necessitation inference rules:

$$\frac{\varphi, \varphi \rightarrow \psi}{\psi} \quad \frac{\varphi}{B_{\varnothing}^{\varnothing} \varphi}.$$

In addition to the unary relation $\vdash \varphi$, we also consider a binary relation $F \vdash \varphi$. We write $F \vdash \varphi$ if a formula $\varphi \in \Phi$ is provable from the set of formulae $F \subseteq \Phi$ and the theorems of our logical system using only the Modus Ponens inference rule. It is easy to see that statement $\varnothing \vdash \varphi$ is equivalent to $\vdash \varphi$. We say that set F is inconsistent if there is a formula $\varphi \in F$ such that $F \vdash \varphi$ and $F \vdash \neg \varphi$.

The proof of the following theorem is in the full version of this paper (Jiang and Naumov 2023).

Theorem 1 (soundness) *If $\vdash \varphi$, then $w \Vdash \varphi$ for each state w of an arbitrary game.*

We conclude this section with two auxiliary results that will be used later in the proof of the completeness. The first of them is a form of the Positive Introspection principle for our belief modality B .

Lemma 1 $\vdash B_X^T \varphi \rightarrow B_X^{\varnothing} B_X^T \varphi$.

PROOF. Formula $B_X^{\varnothing} \neg B_X^T \varphi \rightarrow \neg B_X^T \varphi$ is an instance of the Truth axiom. Thus, $\vdash B_X^T \varphi \rightarrow \neg B_X^{\varnothing} \neg B_X^T \varphi$, by contraposition. Hence, taking into account the following instance $\neg B_X^{\varnothing} \neg B_X^T \varphi \rightarrow B_X^{\varnothing} \neg B_X^{\varnothing} \neg B_X^T \varphi$ of the Negative Introspection axiom, we have

$$\vdash B_X^T \varphi \rightarrow B_X^{\varnothing} \neg B_X^{\varnothing} \neg B_X^T \varphi. \quad (5)$$

Also, the formula $\neg B_X^T \varphi \rightarrow B_X^{\varnothing} \neg B_X^T \varphi$ is an instance of the Negative Introspection axiom. Thus, by contraposition, $\vdash \neg B_X^{\varnothing} \neg B_X^T \varphi \rightarrow B_X^T \varphi$. Hence, by the Necessitation inference rule, $\vdash B_{\varnothing}^{\varnothing} (\neg B_X^{\varnothing} \neg B_X^T \varphi \rightarrow B_X^T \varphi)$. Thus, $\vdash B_X^{\varnothing} (\neg B_X^{\varnothing} \neg B_X^T \varphi \rightarrow B_X^T \varphi)$ by the Monotonicity axiom and the Modus Ponens inference rule. Thus, by the Distributivity axiom and the Modus Ponens inference rule it follows that $\vdash B_X^{\varnothing} \neg B_X^{\varnothing} \neg B_X^T \varphi \rightarrow B_X^{\varnothing} B_X^T \varphi$. The latter, together with statement (5), implies the statement of the lemma by propositional reasoning. \square

The proof of the following lemma is in the full version of this paper (Jiang and Naumov 2023).

Lemma 2 *If $\varphi_1, \dots, \varphi_n \vdash \psi$, then $B_X^T \varphi_1, \dots, B_X^T \varphi_n \vdash B_X^T \psi$.*

5 Strong Completeness

The proof of the completeness theorem is split into four subsections. In Subsection 5.1, we define the canonical game. The truth lemma for the canonical game and the final step of the proof are given in Subsection 5.4. In the two subsections before, we state and prove auxiliary lemmas used in the proof of the truth lemma. A highly non-trivial proof of one of these auxiliary lemmas is located in (Jiang and Naumov 2023).




Figure 2: Fragment of a canonical tree.

5.1 Canonical Game

Throughout the rest of the paper, we assume that \mathcal{B} is any set of cardinality larger than \mathcal{A} , such as, for example, the powerset $\mathcal{P}(\mathcal{A})$. We explain the need for \mathcal{B} after Definition 3 below.

In this section, towards the proof of completeness, for an arbitrary maximal consistent set of formulae $F_0 \subseteq \Phi$ and an arbitrary dataset $T_0 \subseteq V$, we define a ‘‘canonical’’ game $G(T_0, F_0) = (W, \{\sim_x\}_{x \in V}, \{\mathcal{T}_w\}_{w \in W}, \Delta, M, \pi)$.

The set of states in the canonical game is defined using the tree construction which goes back to the proof of the completeness theorem for epistemic logic of distributed knowledge (Fagin, Halpern, and Vardi 1992). Informally, in the tree construction, the states are the nodes of a tree. Formally, the states are defined first as finite sequences; the tree structure on these sequences is specified later.

Definition 3 *W is the set of all sequences $T_0, F_0, X_1, b_1, T_1, F_1, \dots, X_n, b_n, T_n, F_n$ such that $n \geq 0$ and*

1. $F_i \subseteq \Phi$ is a maximal consistent set of formulae for each integer $i \geq 1$,
2. $T_i, X_i \subseteq V$ are datasets for each integer $i \geq 1$,
3. $b_i \in \mathcal{B}$ for each integer $i \geq 1$,
4. $\{\varphi \mid B_{X_i}^{\varnothing} \varphi \in F_{i-1}\} \subseteq F_i$ for each integer $i \geq 1$,
5. $B_Y^T \psi \rightarrow \psi \in F_i$ for each dataset $Y \subseteq V$, each formula $\psi \in \Phi$, and each integer $i \geq 0$.

For any two states $u, w \in W$ of the form

$$u = T_0, F_0, \dots, X_{n-1}, b_{n-1}, T_{n-1}, F_{n-1}$$

$$w = T_0, F_0, \dots, X_{n-1}, b_{n-1}, T_{n-1}, F_{n-1}, X_n, b_n, T_n, F_n,$$

we say that the states are *adjacent*. Note that the adjacency relation forms a tree structure (undirected graph without cycles) on set W . We say that the edge (u, w) is *labelled* with each data variable in set X_n . By $F(w)$ and $T(w)$ we denote sets F_n and T_n , respectively. Informally, it is convenient to visualise this tree with the edge marked by the pair X_n, b_n and node w marked by the pair T_n, F_n . Figure 2 depicts a fragment of such visual representation. In this figure, the node $T_0, F_0, X_1, b_1, T_1, F_1$ is adjacent to the node $T_0, F_0, X_1, b_1, T_1, F_1, X_3, b_3, T_3, F_3$. The edge between them is *labelled* with each variable in set X_3 .

By *clones* we call any two nodes

$$w_1 = T_0, F_0, \dots, X_{n-1}, b_{n-1}, T_{n-1}, F_{n-1}, X_n, b_n, T_n, F_n$$

$$w_2 = T_0, F_0, \dots, X_{n-1}, b_{n-1}, T_{n-1}, F_{n-1}, X_n, b'_n, T_n, F_n$$

that differ only by the *last* b -value. For example, in Figure 2, nodes $T_0, F_0, X_1, b_1, T_1, F_1, X_3, b_3, T_3, F_3$ and $T_0, F_0, X_1, b_1, T_1, F_1, X_3, b_6, T_3, F_3$ are clones. The purpose of set \mathcal{B} in our construction is to guarantee that each node (except for the root) has more clones than the cardinality of set \mathcal{A} .

Recall that a *simple* path in a graph is a path without self-intersections and that in any tree there is a unique simple path between any two nodes.

Definition 4 For any states $w, w' \in W$ and any variable $x \in V$, let $w \sim_x w'$ if every edge along the unique simple path between nodes w and w' is labelled with variable x .

Lemma 3 Relation \sim_x is an equivalence relation on set W for each data variable $x \in V$.

As we prove later in Lemma 9, for any state $w \in W$, set $F(w)$ is the set of all formulae that are satisfied in state w of our canonical model. At the same time, set $T(w)$ is the set of all data variables that are trustworthy in state w :

Definition 5 $\mathcal{T}_w = T(w)$ for any state $w \in W$.

The above definition explains the intuition behind item 5 of Definition 3. Indeed, the item states $B_Y^{T(w)}\varphi \rightarrow \varphi \in F(w)$. Intuitively, this means that if $B_Y^{T(w)}\varphi$ is true in a state w , where set $T(w)$ is trustworthy by Definition 5, then statement φ must be true in state w .

In this paper, as in several other works that extend Marc Pauly's logic of coalitional power (Goranko and van Drimelen 2006; Naumov and Tao 2017; Goranko and Enqvist 2018), the mechanism of the canonical game is using the voting construction. In the standard version of this construction, if the goal of the coalition is to achieve φ , then all members of the coalition use action φ . Informally, this could be interpreted as "voting for φ ". In our case, a strategy is informed by a dataset X . Thus, for the strategy to work the coalition should somehow prove that it has access to the values of dataset X in the current state. To achieve this, we require each vote to be "signed" by naming a state that belongs to the same X -equivalence class as the current state. Thus, each action consists of a formula and a state ("signature"). Similar constructions are used in (Naumov and Tao 2018c; Jiang and Naumov 2022a).

Definition 6 $\Delta = \Phi \times W$.

If d is a pair (x, y) , then by $pr_1(d)$ and $pr_2(d)$ we mean elements x and y , respectively.

Definition 7 Mechanism M consists of all triples $(w, \delta, v) \in W \times \Delta^A \times W$ such that for each formula $[C]_X^T\varphi \in F(w)$, if

1. $pr_1(\delta(a)) = \varphi$ for each actor $a \in C$ and
2. $w \sim_X pr_2(\delta(a))$ for each actor $a \in C$,
3. $T \subseteq \mathcal{T}_w$,
4. $T \subseteq \mathcal{T}_v$,

then $\varphi \in F(v)$.

Definition 8 $\pi(p) = \{w \in W \mid p \in F(w)\}$ for each atomic proposition $p \in P$.

This concludes the definition of the canonical game. As usual, the key step in the proof of completeness is a "truth lemma" proven by induction on the structural complexity of a formula. In our case, this is Lemma 9. To improve the readability of the proof of Lemma 9, we separate the non-trivial induction steps into separate lemmas stated in the next two subsections.

5.2 Properties of the Belief Modality

In this subsection, we state and prove lemmas for modality B used in the induction step of the proof of Lemma 9.

Lemma 4 For any formula $B_Y^T\varphi \in \Phi$ and any states

$$\begin{aligned} w' &= T_0, F_0, \dots, b_{n-1}, T_{n-1}, F_{n-1}, \\ w &= T_0, F_0, \dots, b_{n-1}, T_{n-1}, F_{n-1}, X_n, b_n, T_n, F_n \end{aligned}$$

if $Y \subseteq X_n$, then $B_Y^T\varphi \in F(w')$ iff $B_Y^T\varphi \in F(w)$.

PROOF. (\Rightarrow) : Let $B_Y^T\varphi \in F(w')$. Then, $B_Y^T\varphi \in F_{n-1}$. Thus, by Lemma 1 and the Modus Ponens inference rule, $F_{n-1} \vdash B_Y^{\emptyset}B_Y^T\varphi$. Hence, $F_{n-1} \vdash B_{X_n}^{\emptyset}B_Y^T\varphi$ by the assumption $Y \subseteq X_n$ of the lemma, the Monotonicity axiom, and the Modus Ponens inference rule. Thus, $B_{X_n}^{\emptyset}B_Y^T\varphi \in F_{n-1}$ because F_{n-1} is a maximal consistent set. Then, $B_Y^T\varphi \in F_n$ by item 4 of Definition 3. Therefore, $B_Y^T\varphi \in F(w)$.

(\Leftarrow) : Let $B_Y^T\varphi \notin F(w')$. Thus, $B_Y^T\varphi \notin F_{n-1}$. Hence, $\neg B_Y^T\varphi \in F_{n-1}$ because F_{n-1} is a maximal consistent set of formulae. Then, $F_{n-1} \vdash B_Y^{\emptyset}\neg B_Y^T\varphi$ by the Negative Introspection axiom and the Modus Ponens inference rule. Thus, $F_{n-1} \vdash B_{X_n}^{\emptyset}\neg B_Y^T\varphi$ by the assumption $Y \subseteq X_n$ of the lemma, the Monotonicity axiom, and the Modus Ponens inference rule. Hence, because set F_{n-1} is maximal, $B_{X_n}^{\emptyset}\neg B_Y^T\varphi \in F_{n-1}$. Then, $\neg B_Y^T\varphi \in F_n$ by item 4 of Definition 3. Thus, $B_Y^T\varphi \notin F_n$, because set F_n is consistent. Therefore, $B_Y^T\varphi \notin F(w)$. \square

Lemma 5 For any states $w, u \in W$ and any formula $B_X^T\varphi \in F(w)$, if $w \sim_X u$ and $T \subseteq \mathcal{T}_w$, then $\varphi \in F(u)$.

PROOF. By Definition 4, the assumption $w \sim_X u$ implies that each edge along the unique path between nodes w and u is labelled with each variable in dataset X . Then, the assumption $B_X^T\varphi \in F(w)$ implies $B_X^T\varphi \in F(u)$ by applying Lemma 4 to each edge along this path. Note that the assumption $T \subseteq \mathcal{T}_w$ of the lemma implies that $T \subseteq T(u)$ by Definition 5. Hence, $F(u) \vdash B_X^{T(u)}\varphi$ by the Monotonicity axiom and the Modus Ponens inference rule. Thus, $F(u) \vdash \varphi$ by item 5 of Definition 3 and the Modus Ponens inference rule. Therefore, $\varphi \in F(u)$ because the set $F(u)$ is maximal. \square

Lemma 6 For any $w \in W$ and any formula $B_X^T\varphi \notin F(w)$, there exists a state $u \in W$ such that $w \sim_X u$, $T = \mathcal{T}_w$, and $\varphi \notin F(u)$.

PROOF. Consider the set of formulae

$$\begin{aligned} H &= \{\neg\varphi\} \cup \{\psi \mid B_X^{\emptyset}\psi \in F(w)\} \\ &\quad \cup \{B_Y^T\chi \rightarrow \chi \mid Y \subseteq V, \chi \in \Phi\}. \end{aligned} \quad (6)$$

Claim 1 Set H is consistent.

PROOF OF CLAIM. Assume the opposite. Hence, there are formulae $\chi_1, \dots, \chi_n \in \Phi$, datasets $Y_1, \dots, Y_n \subseteq V$, and formulae

$$B_X^\varnothing \psi_1, \dots, B_X^\varnothing \psi_m \in F(w) \quad (7)$$

such that

$$B_{Y_1}^T \chi_1 \rightarrow \chi_1, \dots, B_{Y_n}^T \chi_n \rightarrow \chi_n, \psi_1, \dots, \psi_m \vdash \varphi.$$

Thus, by Lemma 2,

$$B_X^T (B_{Y_1}^T \chi_1 \rightarrow \chi_1), \dots, B_X^T (B_{Y_n}^T \chi_n \rightarrow \chi_n), \\ B_X^T \psi_1, \dots, B_X^T \psi_m \vdash B_X^T \varphi.$$

Hence, $B_X^T \psi_1, \dots, B_X^T \psi_m \vdash B_X^T \varphi$ by the Trust axiom applied n times. Then, $B_X^\varnothing \psi_1, \dots, B_X^\varnothing \psi_m \vdash B_X^T \varphi$ by the Monotonicity axiom and the Modus Ponens inference rule applied m times. Thus, $F(w) \vdash B_X^T \varphi$ due to statement (7). Hence, $B_X^T \varphi \in F(w)$ because the set $F(w)$ is maximal, which contradicts the assumption $B_X^T \varphi \notin F(w)$ of the lemma. \square

Let H' be any maximal consistent extension of set H and b be any element of set \mathcal{B} . Suppose that $w = T_0, F_0, \dots, X_n, b_n, T_n, F_n$. Consider sequence

$$u = T_0, F_0, \dots, X_n, b_n, T_n, F_n, X, b, T, H'. \quad (8)$$

Note that $u \in W$ by Definition 3, equation (6), and the choice of set H' as an extension of set H . Also, observe that $w \sim_X u$ by Definition 4 and equation (8). In addition, $T = T(u) = \mathcal{T}_u$ by equation (8) and Definition 5. Finally, $\neg\varphi \in H \subseteq H' = F(u)$ by equation (6), the choice of H' as an extension of H , and equation (8). Therefore, $\varphi \notin F(u)$ because the set $F(u)$ is consistent. This concludes the proof of the lemma. \square

5.3 Properties of the Doxastic Strategy Modality

We now state lemmas for modality $[C]$ that will be used in the induction step of the proof of Lemma 9.

Lemma 7 *For an arbitrary state $w \in W$ and any formula $[C]_X^T \varphi \in F(w)$ there is an action profile $s \in \Delta^C$ such that for all states $w', v \in W$ and each complete action profile $\delta \in \Delta^A$ if $w \sim_X w'$, $s =_C \delta$, $T \subseteq \mathcal{T}_{w'}$, $T \subseteq \mathcal{T}_v$, and $(w', \delta, v) \in M$, then $\varphi \in F(v)$.*

PROOF. Let action profile $s \in \Delta^C$ be such that

$$s(a) = (\varphi, w) \quad (9)$$

for each actor $a \in C$. Consider any states $w', v \in W$ and any complete action profile $\delta \in \Delta^A$ such that

$$w \sim_X w', \quad (10)$$

$$s =_C \delta, T \subseteq \mathcal{T}_{w'}, T \subseteq \mathcal{T}_v, \text{ and } (w', \delta, v) \in M. \quad (11)$$

It suffices to show that $\varphi \in F(v)$.

The assumption $[C]_X^T \varphi \in F(w)$ of the lemma implies $F(w) \vdash B_X^T [C]_X^T \varphi$ by the Strategic Introspection axiom and propositional reasoning. Hence, $B_X^T [C]_X^T \varphi \in F(w)$ because set $F(w)$ is maximal. Thus, $[C]_X^T \varphi \in F(w')$ by Lemma 5 and assumption (10) and the part $T \subseteq \mathcal{T}_{w'}$ of assumption (11). Therefore, $\varphi \in F(v)$ by Definition 7, and assumptions (9) and (11). \square

The proof of the next lemma can be found in (Jiang and Naumov 2023).

Lemma 8 *For an arbitrary state $w \in W$, any formula $\neg[C]_X^T \varphi \in F(w)$, and any action profile $s \in \Delta^C$, there are states $w', v \in W$ and a complete action profile $\delta \in \Delta^A$ such that $w \sim_X w'$, $s =_C \delta$, $T \subseteq \mathcal{T}_{w'}$, $T \subseteq \mathcal{T}_v$, $(w', \delta, v) \in M$, and $\varphi \notin F(v)$.*

5.4 Final Steps

We are now ready to prove the truth lemma for our logical system and the strong completeness of the system.

Lemma 9 *$w \Vdash \varphi$ iff $\varphi \in F(w)$ for each state $w \in W$ and each formula $\varphi \in \Phi$.*

PROOF. We prove the statement by induction on the complexity of formula φ . Suppose that formula φ is an atomic proposition p . Note that $w \Vdash p$ iff $w \in \pi(p)$ by item 1 of Definition 2. At the same time, $w \in \pi(p)$ iff $p \in F(w)$ by Definition 8. Therefore, $w \Vdash p$ iff $p \in F(w)$.

If formula φ is a negation or an implication, then the statement of the lemma follows from the maximality and the consistency of the set $F(w)$, items 2 and 3 of Definition 2, and the induction hypothesis in the standard way.

Suppose that formula φ has the form $B_X^T \psi$.

(\Rightarrow) : Assume that $B_X^T \psi \notin F(w)$. Then, $\neg B_X^T \psi \in F(w)$ because $F(w)$ is a maximal consistent set of formulae. Thus, by Lemma 6, there is a state $w' \in W$ such that $w \sim_X w'$, $T \subseteq \mathcal{T}_{w'}$, and $\neg\psi \in F(w')$. Hence, $\psi \notin F(w')$ because set $F(w')$ is consistent. Then, $w' \not\Vdash \psi$ by the induction hypothesis. Therefore, $w \not\Vdash B_X^T \psi$ by item 4 of Definition 2 and statements $w \sim_X w'$ and $T \subseteq \mathcal{T}_{w'}$.

(\Leftarrow) : Assume that $B_X^T \psi \in F(w)$. Consider any state w' such that $w \sim_X w'$ and $T \subseteq \mathcal{T}_{w'}$. By item 4 of Definition 2, it suffices to show that $w' \Vdash \psi$, which is true by Lemma 5.

Finally, suppose that formula φ has the form $[C]_X^T \psi$.

(\Rightarrow) : Assume that $[C]_X^T \psi \notin F(w)$. Thus, $\neg[C]_X^T \psi \in F(w)$ because set $F(w)$ is maximal. Hence, by Lemma 8, for any action profile $s \in \Delta^C$, there are states $w', v \in W$ and a complete action profile $\delta \in \Delta^A$ such that $w \sim_X w'$, $s =_C \delta$, $T \subseteq \mathcal{T}_{w'}$, $T \subseteq \mathcal{T}_v$, $(w', \delta, v) \in M$, and $\psi \notin F(v)$. Then, by the induction hypothesis, for any action profile $s \in \Delta^C$, there are states $w', v \in W$ and a complete action profile $\delta \in \Delta^A$ such that $w \sim_X w'$, $s =_C \delta$, $T \subseteq \mathcal{T}_{w'}$, $T \subseteq \mathcal{T}_v$, $(w', \delta, v) \in M$, and $v \not\Vdash \psi$. Therefore, $w \not\Vdash [C]_X^T \psi$ by item 5 of Definition 2.

(\Leftarrow) : Assume $[C]_X^T \psi \in F(w)$. Thus, by Lemma 7, there is an action profile $s \in \Delta^C$ such that for all states $w', v \in W$ and each complete action profile $\delta \in \Delta^A$ if $w \sim_X w'$, $s =_C \delta$, $T \subseteq \mathcal{T}_{w'}$, $T \subseteq \mathcal{T}_v$, and $(w', \delta, v) \in M$, then $\psi \in F(v)$. Hence, by the induction hypothesis, there is an action profile $s \in \Delta^C$ such that for all states $w', v \in W$ and each complete action profile $\delta \in \Delta^A$ if $w \sim_X w'$, $s =_C \delta$, $T \subseteq \mathcal{T}_{w'}$, $T \subseteq \mathcal{T}_v$, and $(w', \delta, v) \in M$, then $v \Vdash \psi$. Therefore, $w \Vdash [C]_X^T \psi$ by item 5 of Definition 2. \square

Theorem 2 (strong completeness) *For any set of formulae $F \subseteq \Phi$ and any formula $\varphi \in \Phi$, if $F \not\Vdash \varphi$, then there is a*

state w of a game such that $w \Vdash f$ for each formula $f \in F$ and $w \not\Vdash \varphi$.

PROOF. The assumption $F \not\vdash \varphi$ implies that the set $F \cup \{\neg\varphi\}$ is consistent. Let F_0 be any maximal consistent extension of this set. Consider the canonical game $G(\emptyset, F_0)$.

First, we show that the sequence \emptyset, F_0 is a state of this canonical game. By Definition 3, it suffices to show that $B_Y^\emptyset \psi \rightarrow \psi \in F_0$ for each dataset $Y \subseteq V$ and each formula $\psi \in \Phi$. The last statement is true by the Truth axiom and because set F_0 is maximal.

Finally, note that $\varphi \notin F_0$ because set F_0 is consistent and $\neg\varphi \in F_0$. Therefore, by Lemma 9 and because $F \subseteq F_0$, it follows that $\emptyset, F_0 \Vdash f$ for each formula $f \in F$ and also $\emptyset, F_0 \not\Vdash \varphi$. \square

6 Ex Ante and Ex Post Trust

In item 5 of Definition 2, we require that $T \subseteq \mathcal{T}_u$ and $T \subseteq \mathcal{T}_v$. In other words, we apply the assumption of trustworthiness of dataset T *ex ante* (before action) and *ex post* (after action). In general, one can consider that different datasets, A and P , are required to be trustworthy *ex ante* and *ex post*, respectively. That leads to a more general modality $[C]_X^{A,P} \varphi$, defined below:

$w \Vdash [C]_X^{A,P} \varphi$, when there is an action profile $s \in \Delta^C$ of coalition C such that for all states $u, v \in W$ and each complete action profile $\delta \in \Delta^A$ if $w \sim_X u$, $A \subseteq \mathcal{T}_u$, $s =_C \delta$, $(u, \delta, v) \in M$, and $P \subseteq \mathcal{T}_v$, then $v \Vdash \varphi$.

One might wonder which of the data variables among \vec{x} , \vec{v} , and t should be trusted *ex ante* and which *ex post* for statement (4) to be true in our introductory example. Unfortunately, we do not have access to the Patriot code to answer this question, but we suspect that the Patriot missile constantly adjusts the trajectory based on the current speed and position of the target. To model such behaviour one would need to use multi-step games instead of one-shot (strategic) games that we consider in this paper.

To illustrate *ex ante* and *ex post* trust, let us consider a different example where a governing body consisting of 25 members is about to vote on passing a certain regulation. Suppose that each of them votes *yes* or *no* by a paper ballot. After the vote, the ballots are counted by a tallyman and the number of *yes* votes, denoted by n , is announced. If n is more than 12, then the regulation is approved, see Figure 3.




Figure 3: Voting game.

Let us now further assume that a newspaper contacts all 25 members and ask them how they plan to vote. Suppose that 10 members state that they plan to support the proposal, 10 members say that they plan to vote against the proposal,

and 5 members are undecided. Alice, Bob, and Cathy are among those 5 who are undecided. It appears that the coalition consisting of the three of them has a strategy to sway the outcome of the vote either way. If all three of them vote *yes*, then the regulation is approved. If they vote *no*, then it is rejected. Note, however, that both of these strategies are *doxastic*: they rely on the *ex ante* trust in the newspaper data and on the *ex post* trust in the tallyman computing n correctly. Let data variables yea and nay represent the number of members, according to the newspaper, who plan to vote for and against, respectively. Note that, in the current world w , $yea = nay = 10$. Thus,

$$\begin{aligned} w \Vdash [Alice, Bob, Cathy]_{\{yea\}}^{\{yea\}, \{n\}} \text{“Approved”}, \\ w \Vdash [Alice, Bob, Cathy]_{\{nay\}}^{\{nay\}, \{n\}} \text{“Rejected”}. \end{aligned} \quad (12)$$

All subscripts and superscripts are important in both formulae above. For example, if subscript yea is dropped in statement (12), then the statement is no longer true:

$$w \not\Vdash [Alice, Bob, Cathy]_{\emptyset}^{\{yea\}, \{n\}} \text{“Approved”}.$$

This is because, although the doxastic strategy exists, the empty dataset *does not inform* its existence.

If yea is removed from the superscript, then the statement is also not true

$$w \not\Vdash [Alice, Bob, Cathy]_{\{yea\}}^{\emptyset, \{n\}} \text{“Approved”}$$

because the newspaper’s information is no longer assumed to be trusted. Without the *ex ante* trust in the existence of at least 10 members who are ready to support the regulation, the coalition consisting of Alice, Bob, and Cathy does not have a strategy to pass the regulation.

Finally, without the *ex post* trust that the votes will be counted correctly, such strategy does not exist either:

$$w \not\Vdash [Alice, Bob, Cathy]_{\{yea\}}^{\{yea\}, \emptyset} \text{“Approved”}.$$

Note that formula $[C]_X^{T,T} \varphi$ is equivalent to our original modality $[C]_X^T \varphi$.

The axioms of our logical system could be modified in a straightforward way for this more general modality. The most significant change is for the Public Belief axiom: $B_{\emptyset}^T \varphi \rightarrow [\emptyset]_{\emptyset}^{\emptyset, T} \varphi$. Although we did not check the details, we believe that our soundness and completeness results can be adjusted for this more general modality.

7 Conclusion

In this paper, we proposed the notion of doxastic strategies. We gave a sound and complete logical system for reasoning about such strategies and outlined a possible extension of this work that separates *ex ante* and *ex post* trust.

Acknowledgments

Junli Jiang acknowledges the support of the Innovation Research 2035 Pilot Plan of Southwest University (NO. SWUPilotPlan018).

References

- Ågotnes, T.; and Alechina, N. 2019. Coalition Logic with Individual, Distributed and Common Knowledge. *Journal of Logic and Computation*, 29: 1041–1069.
- Apple, R. W., Jr. 1991. War in the Gulf: Scud Attack; Scud Missile Hits a U.S. Barracks, Killing 27. *New York Times*, February 26th. <https://www.nytimes.com/1991/02/26/world/war-in-the-gulf-scrud-attack-scrud-missile-hits-a-us-barracks-killing-27.html>.
- Baltag, A.; and van Benthem, J. 2021. A simple logic of functional dependence. *Journal of Philosophical Logic*, 50: 1–67.
- Bush, G. H. 1991. Remarks to Raytheon Missile Systems Plant Employees in Andover, Massachusetts. <https://bush41library.tamu.edu/archives/public-papers/2711>. Accessed: 2022-08-20.
- Deuser, K.; Jiang, J.; Naumov, P.; and Zhang, W. 2024. A Dynamic Logic of Data-Informed Knowledge. *Journal of Philosophical Logic*. (to appear).
- Fagin, R.; Halpern, J. Y.; and Vardi, M. Y. 1992. What can machines know? On the properties of knowledge in distributed systems. *Journal of the ACM (JACM)*, 39(2): 328–376.
- Fervari, R.; Herzig, A.; Li, Y.; and Wang, Y. 2017. Strategically knowing how. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, 1031–1038.
- Fervari, R.; Velázquez-Quesada, F. R.; and Wang, Y. 2021. Bisimulations for Knowing How Logics. *The Review of Symbolic Logic*, 1–37.
- Goranko, V.; and Enqvist, S. 2018. Socially Friendly and Group Protecting Coalition Logics. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems*, 372–380. International Foundation for Autonomous Agents and Multiagent Systems.
- Goranko, V.; and van Drimmelen, G. 2006. Complete axiomatization and decidability of Alternating-time temporal logic. *Theoretical Computer Science*, 353(1): 93 – 117.
- Grossi, D.; Lorini, E.; and Schwarzenrüber, F. 2015. The ceteris paribus structure of logics of game forms. *Journal of Artificial Intelligence Research*, 53: 91–126.
- Jiang, J.; and Naumov, P. 2022a. Data-Informed Knowledge and Strategies. *Artificial Intelligence*, 309: 103727.
- Jiang, J.; and Naumov, P. 2022b. In Data We Trust: The Logic of Trust-Based Beliefs. In *the 31st International Joint Conference on Artificial Intelligence (IJCAI-22)*.
- Jiang, J.; and Naumov, P. 2023. The Logic of Doxastic Strategies. *arXiv:2312.07107*.
- Legion, T. A. 2017. 14th Quartermaster Detachment Memorial. <https://www.legion.org/memorials/238236/14th-quartermaster-detachment-memorial>. Accessed: 2022-08-20.
- Li, Y.; and Wang, Y. 2021. Planning-based knowing how: A unified approach. *Artificial Intelligence*, 103487.
- Mignotte, M. 2010. The Patriot Missile Failure. <https://www.iro.umontreal.ca/~mignotte/IFT2425/Disasters.html>. Accessed: 2022-08-20.
- Naumov, P.; and Tao, J. 2017. Coalition Power in Epistemic Transition Systems. In *Proceedings of the 2017 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 723–731.
- Naumov, P.; and Tao, J. 2018a. Second-Order Know-How Strategies. In *Proceedings of the 2018 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 390–398.
- Naumov, P.; and Tao, J. 2018b. Strategic Coalitions with Perfect Recall. In *Proceedings of Thirty-Second AAAI Conference on Artificial Intelligence*.
- Naumov, P.; and Tao, J. 2018c. Together We Know How to Achieve: An Epistemic Logic of Know-How. *Artificial Intelligence*, 262: 279 – 300.
- Pauly, M. 2002. A Modal Logic for Coalitional Power in Games. *Journal of Logic and Computation*, 12(1): 149–166.