

Invariance Properties of Quantifiers and Multiagent Information Exchange

Nina Gierasimczuk* and Jakub Szymanik†

Institute of Artificial Intelligence, University of Groningen
nina.gierasimczuk@gmail.com
jakub.szymanik@gmail.com

Abstract. The paper presents two case studies of multi-agent information exchange involving generalized quantifiers. We focus on scenarios in which agents successfully converge to knowledge on the basis of the information about the knowledge of others, so-called Muddy Children puzzle [1] and Top Hat puzzle. We investigate the relationship between certain invariance properties of quantifiers and the successful convergence to knowledge in such situations. We generalize the scenarios to account for public announcements with arbitrary quantifiers. We show that the Muddy Children puzzle is solvable for any number of agents if and only if the quantifier in the announcement is positively active (satisfies a version of the variety condition). In order to get the characterization result, we propose a new concise logical modeling of the puzzle based on the number triangle representation of generalized quantifiers. In a similar vein, we also study the Top Hat puzzle. We observe that in this case an announcement needs to satisfy stronger conditions in order to guarantee solvability. Hence, we introduce a new property, called bounded thickness, and show that the solvability of the Top Hat puzzle for arbitrary number of agents is equivalent to the announcement being 1-thick.

Key words: generalized quantifiers; number triangle; invariance properties; Muddy Children Puzzle; Top Hat Puzzle, epistemic logic

1 Introduction

The Top Hat puzzle Imagine you are one of ten prisoners locked up for extensive use of logic. To make you even more miserable, the guard comes up with a puzzle. He gathers all ten of you and says: ‘Each of you will be assigned a random hat, either black or white. You will be lined up single file where each can see the hats in front of him but not behind. Starting with the prisoner in the back of the line and moving forward, you must each, in turn, say only one word which must be ‘black’ or ‘white’. If the word you uttered matches your hat color you are released, if not, you are killed on the spot. You have half an hour to pray for your life.’ Then he leaves. One of the prisoners says: ‘I have a plan! If you

*The author’s research was supported by Vidi Grant NWO-639.072.904.

†The author’s research was supported by Vici Grant NWO-277-80-001.

agree on it, 9 of us 10 will definitely survive, and the remaining one has a 50/50 chance of survival.’ What does he have in mind?

Most probably the strategy that he wants to implement is as follows. First, the prisoners have to agree on the following meaning of the utterance of the one who is the last in the line. If he says ‘white’, it means that he sees an even number of white hats in front of him. If he says ‘black’ it means that he sees an odd number of white hats in front of him. Hence, his utterance has nothing to do with what he thinks his own hat is—he simply announces the parity of black hats among the remaining 9 prisoners. There is a 50/50 chance of the total number black hats being odd or even, and a 50/50 chance of his hat being black or white, and the same is his chance of survival. However, after this utterance the prisoner that stands in front of him knows for sure the color of his hat—he compares the utterance of his predecessor with the number of white hats he sees in front of him. If the parity is the same, he concludes that his hat is black, otherwise it is white. He announces his conclusion. Now the person in front of him takes into account the first announcement and the second utterance, sees the number of white hats in front of her, and now she is also certain about her hat’s color, etc.

The Muddy Children Puzzle Yet another thought experiment—you are now out of prison, visiting a relative, who has three children. While you are having coffee in the living-room, the kids are playing outside. When they come back home, their father says: (1) ‘At least one of you has mud on your forehead’. Then, he asks the children: (I) ‘Can you tell for sure whether you have mud on your forehead? If yes, announce your status’. Children know that their father never lies and that they are all perfect logical reasoners. Each child can see the mud on others but cannot see his or her own forehead. Nothing happens. But after the father repeats the question for the second time suddenly all muddy children know that they have mud on their forehead. How is that possible?

Let us again explain away the surprising outcome, this time using a formal machinery. The problem can be modeled with the help of Kripke structures describing agents’ uncertainty. Let us give the three children names: a , b and c , and assume that, in fact, all of them are muddy. Three propositional letters m_a , m_b and m_c express that the corresponding child is muddy. The standard epistemic modeling is depicted in Figure 1, with the initial model of the situation on the left (see [2]).

In the model, possible worlds correspond to the ‘distribution of mud’ on children’s foreheads, e.g., $w_5 : m_a$ stands for a being muddy and b and c being clean in world w_5 . Two worlds are joined with an edge labelled with x , if they are in the uncertainty range of agent x (i.e., if agent x cannot distinguish between the two worlds; for clarity we drop the reflexive arrows for each state). The boxed state stands for the actual world. Now, let us recall how the solution process can be modeled in this setting. The first public announcement has the following form: (1’) $m_a \vee m_b \vee m_c$, and after the announcement (1’) becomes common knowledge among children. As a result the children perform an update, i.e., they eliminate world w_8 in which (1’) is false. The result is depicted in the second part of Figure

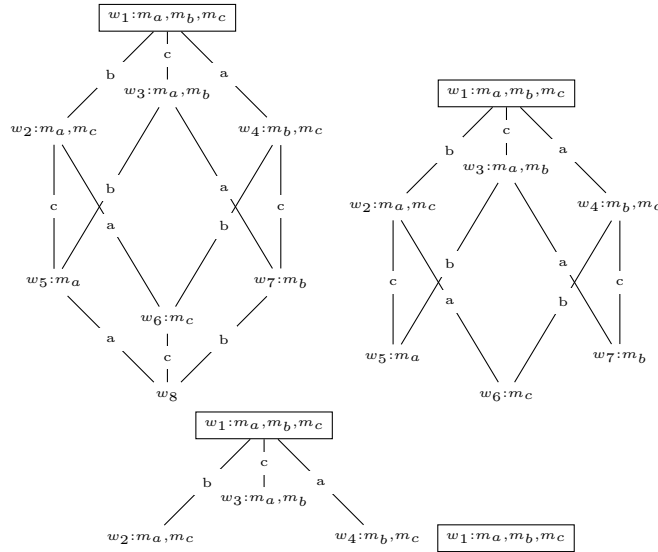


Fig. 1. The Classical Muddy Children Modeling

1. Then the father asks for the first time, who of them knows his status (**I**). The agents' reasoning can be as follows. In world w_6 agent c knows that he is dirty (there is no uncertainty of agent c between this world and another in which he is clean). Therefore, if the actual world was w_6 , agent c would know his state and announce it. The situation is similar for a and b in w_5 and w_7 , respectively. The silence of the children may also be interpreted as the announcement that none of them know whether they are muddy. Hence, all agents eliminate those worlds that do not make such announcement true: w_5, w_6, w_7 . The epistemic model of the next stage is smaller by three worlds. Then it is again clear that if one of the w_2, w_3 , or w_4 was the actual state, the respective agents would have announced their knowledge. The children still do not respond so, in the next round, everyone knows that the actual situation cannot be any of w_2, w_3 , and w_4 . Hence, they all eliminate the three states, which leaves them with just one possibility, w_1 . All uncertainty disappears and they all know that they are dirty at the same time.¹

The above epistemic scenarios show that a very simple *quantitative* public announcement carries powerful *qualitative* information. The similarity between the Muddy Children puzzle and the Top Hats problem is striking: in both cases agents need to reason about their properties on the basis of some general quantitative statement; the settings differ with respect to the observational power of the agents. Intuitively, the possibility of convergence to knowledge in such problems depends on the trade-off between the internal structure of epistemic information and the amount of information provided by the public announcement. To see these differences in full light let us consider the following two cases:

¹Note that the reasoning in the Top Hat puzzle can be modeled in a similar way.

- The Top Hats puzzle: announcing ‘an even number of hats are white’ allows epistemic reasoning that solves the puzzle for any configuration; announcing ‘at least one hat is black’ allows solving the problem only in a very limited number of cases.
- The Muddy Children puzzle: announcing ‘at least one of you has mud on your forehead’ allows epistemic reasoning that solves the puzzle for any configuration, while announcing parity leads to an immediate one-step solution that does not involve any epistemic reasoning.

Hence, it is fair to say that in some sense parity announcements bring more information than existential announcements, at least with respect to the above-mentioned epistemic situations.

2 Generalized Muddy Children Puzzle²

Let us recall the father’s first announcement in the Muddy Children puzzle. It has the following form: (1) ‘At least one of you has mud on your forehead’. Sentence (1) can be seen as a *background assumption* that makes the epistemic multi-agent inferential process possible. In a way, the quantifier announcement prepares the ground for epistemic reasoning, and enforces a particular structure on the situation, that triggers the successful reasoning. What makes an announcement ‘good’ in this context?

A simple but crucial observation is that the information provided by the father has the following form:

Q of you have mud on your forehead,

where Q may be substituted by various quantifiers, like ‘At least one’, ‘An even number’, ‘One third’ and so on. Let us think of the Muddy Children situation as $M = (U, A)$, where U is the set of children and $A \subseteq U$ is the set of children that are muddy. Of course, after father’s announcement some models are no longer possible. Only those satisfying the quantifier sentence, i.e., $M \models Q_U(A)$, should be still considered. Therefore, the model of a given Muddy Children scenario consists of the structures satisfying the quantifier sentence. The agent’s goal is to pinpoint one of them—the actual world. To explain this idea in more detail let us start with introducing the notion of generalized quantifiers.

Definition 1 ([5]). *A generalized quantifier Q of type (1) is a class of structures of the form $M = (U, A)$, where A is a subset of U . Additionally, Q is closed under isomorphism, i.e., if M and M' are isomorphic, then $(M \in Q \iff M' \in Q)$.*

Now, the classical Muddy Children puzzle with the father saying ‘At least one of you has mud on your forehead’ involves the existential generalized quantifier:

²This section is an extended discussion of the results published in [3] and reported on in [4].

2.1 Number Triangle Based Modeling of Muddy Children

How can Number Triangle be used to characterize the successful convergence in the generalized Muddy Children puzzle? To answer that question let us analyze a concrete Muddy Children scenario. As before, we take agents a , b , and c . All possibilities with respect to the *size* of the set of muddy children are enumerated in the third level of the number triangle. Let us also assume at this point that the actual situation is that agents a, b are muddy and c is clean. Therefore, with respect to our representation the real world is $(1, 2)$, one child is clean and two are muddy:

$$(3,0) \quad (2,1) \quad \boxed{(1,2)} \quad (0,3)$$

Now, let us focus on what the agents observe. Agent a sees one muddy child and one clean child. The same holds for agent b , in this sense they are perfectly symmetric. Their observational state can be encoded as $(1,1)$. Accordingly, the observational state of c is $(0,2)$. In general, if the number of agents is n , each agent can observe $n - 1$ agents. As a result what agents observe is encoded in the second level of the number triangle.

$$(2,0) \quad (1,1) \quad (0,2)$$

$$(3,0) \quad (2,1) \quad \boxed{(1,2)} \quad (0,3)$$

The question that each of the agents is facing is whether he is muddy. For example, agent a has to decide whether he should *extend* his observation state, $(1, 1)$, to the left state $(2, 1)$ (a decides that he is clean) or to the right state $(1, 2)$ (a decides that he is muddy). The same holds for agent b . The situation of agent c is similar, his observational state is $(0, 2)$ and it has two potential extensions $(1, 2)$ and $(0, 3)$. In general, note that every observational state has two possible successors.

Given this representation, we can now analyze what happens in the Muddy Children scenario. Figure 3 represents the process, with the initial model at the top. First, the announcement is given: ‘At least one of you is muddy’. According to the number triangle representation (see Figure 2 on the right), this allows eliminating those factual states representing finite models that are not in the quantifier. In this case it is $(3, 0)$. The resulting model is the second from the top. Then the father asks: ‘Can you tell for sure whether or not you have mud on your forehead?’ In our graph, this question means: ‘Does any of you have only one successor?’ All agents know that $(3, 0)$ has just been eliminated. Agent a considers it possible that the actual state is $(2, 1)$, i.e., that two agents are clean and one is muddy, so that he himself would have to be clean. But then he knows that there would have to be an agent whose observational state is $(2, 0)$ —there has to be a muddy agent that observes two clean ones. For this hypothetical




Fig. 3. The Number-Triangle Muddy Children Modeling

agent the uncertainty disappeared just after the quantifier announcement (for $(2,0)$ there is only one successor left). So, when it becomes clear that no one knows and the father asks the question again, the world $(2,1)$ gets eliminated and the only possibility for agent a is now $(1,2)$ via the right successor, and this indicates that he has to be muddy. Agent b is in exactly the same situation. They both can announce that they know. And since c witnessed the whole process he knows that the only way for them to know was to be in $(1,1)$ and decides on $(1,2)$.

This epistemic reasoning took two steps. If the actual world was $(2,1)$ some agent's observation would be $(2,0)$, and this agent would know his status after the first announcement, and the rest of the agents would follow. Accordingly, for $(0,3)$ this would have taken three steps. This can be summed up in the following way: the quantifier breaks the perfect 'uncertainty structure' of the model, and the farther the actual state is from this break, the longer it takes to solve the puzzle (as will become clear in Section 2.2).

In general, if there are n agents, we take the n th level of the triangle, i.e., finite models with $|U| = n$, enumerating all possible settings (up to isomorphism). This level will be called the *factual level* and it always consists of $n + 1$ states. It is an analogue of the initial uncertainty domain of the children in the classical modeling. Moreover, in the puzzle every child sees all other children, but not himself, so every possible observation consists of $n - 1$ children. Therefore, level $n - 1$ of the number triangle can be interpreted as enumerating every possible observation of the agents. We will call it the *observational level*. Each observation can be extended to one of the two factual states that are the closest below—to the left if the observer in question is clean or to the right if he is muddy.

2.2 Number of epistemic iterations

By reinterpreting the Muddy Children puzzle within the semantics of quantifiers we can associate every finite model with the number of steps needed to solve the puzzle, if it is solvable at all.

Definition 3. An epistemic quantifier is a pair $Q^{MC} = (Q, f_Q)$, where Q is a quantifier and $f_Q : Q \rightarrow \mathbb{N}$ is a function that assigns to a pair of numbers representing $M \in Q$ the number of steps needed to solve the Muddy Children puzzle with the background assumption containing quantifier Q .

Now, we need to know how to determine values of f_Q for a given quantifier.

Proposition 2. Let Q be a generalized quantifier, and n be the number of children. Then the corresponding epistemic quantifier $Q^{MC} = (Q, f_Q)$, where the partial function $f_Q : Q \rightarrow \mathbb{N}$ is defined in the following way.

$$f_Q((n-m, m)) = \min(\mu_{x \leq n-m} (n-m-x, m+x) \notin Q, \mu_{y \leq m} (n-m+y, m-y) \notin Q).$$

Proof. Observe that the function assigns a value x to $(u-k, k)$ in the level u of the number triangle if $(u-k, k) \in Q$ and there is $(u-\ell, \ell)$ in the level u such that $(u-\ell, \ell) \notin Q$. Moreover, the value x encodes the distance from the nearest $(u-\ell, \ell)$ such that $(u-\ell, \ell) \notin Q$.

Concerning the assignment of the number of steps needed for solving the puzzle, we can also ask what is the structure of those steps. Namely, we can characterize situations in which some agents infer their status from the announcements of other agents, in contrast to the cases in which it happens simultaneously (we use ‘+’-superscripts to identify those situations). The definition of the partial function $f_Q^+ : Q \rightarrow \{+\}$ can be then given in the following way. $f_Q^+((n-m, m)) = +$ iff:

- (1) $f_Q((n-m, m))$ is defined, and
- (2) $m \neq 0$ and $m \neq n$ and some agent considers two factual worlds possible.

For shaping the intuitions, let us give a few examples of epistemic quantifiers in the number triangle representation. First let us consider the quantifier ‘At least k ’. It is easy to observe that increasing k causes the downward triangle to move down along the $(0, 0)$ – $(0, n)$ axis.

This quantifier allows solving the Muddy Children puzzle for any configuration of ‘muddiness’. However, within a certain level, the farther from a minus the longer it takes.

Now let us have a look at the quantifier ‘At most k ’. In Figure 5 the question-marks occur in place of models that satisfy the quantifier, but for which it is impossible to solve the Muddy Children puzzle. For example, if one child is clean and one child is muddy (the actual world is $(1, 1)$) the Muddy Children situation does not lead to a solution if the announcement is: ‘At most two of you are muddy’. Again, the farther from a minus the longer it takes to solve the puzzle.

Parity quantifiers in the Muddy Children setting do not involve much inference—every situation is solvable in one step (see Figure 6) and all answers are given simultaneously by all the agents.



Fig. 4. Increasing muddy-quantifiers ‘At least 1’ and ‘At least 2’




Fig. 5. Decreasing muddy-quantifiers ‘At most 1’ and ‘At most 2’

2.3 Characterization

The above discussion leads to the observation that solving the Muddy Children puzzle is possible if the announcement of the quantifier leaves one observational state with just one successor. Therefore the solvability of the particular Muddy Children scenario can be characterized in the following way:

Theorem 1 (Muddy Children Solvability). *Let n be the number of children, $m \leq n$ the number of muddy children, and Q be the background assumption. A Muddy Children situation is solvable iff $(n - m, m) \in Q$ and there is an $\ell \leq n$ such that $(n - \ell, \ell) \notin Q$.*

Proof. Let us fix n —the number of children and $m \leq n$ —the number of muddy children, Q is the quantifier background assumption.

For left to right. Assume that the scenario ends successfully—all agents arrive to knowledge about their status. Assume towards contradiction that it is not the case that $(n - m, m) \in Q$ or it is not the case that there is $\ell \leq n$ such that $(n - \ell, \ell) \notin Q$.

- if $(n - m, m) \notin Q$ then the father’s announcement is not truthful. Contradiction.
- if for all $\ell \leq n$ it is the case that $(n - \ell, \ell) \in Q$, then the public announcement of Q does not eliminate any world and thus the iterated epistemic reasoning is impossible and the convergence to knowledge fails for all the agents. Contradiction.



Fig. 6. Muddy-quantifiers ‘Divisible by 2’ and ‘Divisible by 3’

For the other direction, assume that $(n - m, m) \in \mathbf{Q}$ and there is $\ell \leq n$ such that $(m - \ell, \ell) \notin \mathbf{Q}$. Then by Proposition 2 $f_Q((n - m, m))$ is defined and hence the puzzle is solvable in $f_Q((n - m, m))$ steps.

In fact, the solvability issue coincides with a known and important property of generalized quantifiers.

Definition 4 ([8]). A quantifier Q is active (alternatively: Q satisfies variety, VAR) iff for every non-empty set U , there exists $A \subseteq U$ such that $Q_U(A)$ but there is also $B \subseteq U$ such that it is not the case that $Q_U(B)$.

Note that VAR can be viewed as a conjunction of two weaker conditions⁴, VAR⁺ and VAR⁻.

Definition 5.

VAR⁺ A quantifier Q is positively active (alternatively: Q satisfies VAR⁺) iff for every non-empty set U if there exists $A \subseteq U$ such that $Q_U(A)$, then there is also $B \subseteq U$ such that it is not the case that $Q_U(B)$.

VAR⁻ A quantifier Q is negatively active (alternatively: Q satisfies VAR⁻) iff for every non-empty set U if there exists $A \subseteq U$ such that it is not the case that $Q_U(A)$, then there is also $B \subseteq U$ such that $Q_U(B)$.

Now, we can characterize the general Muddy Children Solvability in the following way:

Corollary 1 (Muddy Children Solvability) A Muddy Children situation with Q as the background assumption is solvable for any number of children and any distribution of muddiness iff Q is positively active.

2.4 Various Quantifiers as Public Announcements

In this section let us consider a few examples—we will describe Muddy Children quantifiers corresponding to various classes of generalized quantifiers.

⁴Our focus on such forms of VAR is consistent with the usefulness of weaker variability assumptions in Generalized Quantifier Theory [8].

Increasing Quantifiers Let us first consider a class of quantifiers that is closest to the classical version of the Muddy Children puzzle, namely: ‘At least m ’, where $m \in \mathbb{N}$. In the puzzle this takes the form of the announcement: At least k of you have mud on your foreheads. Quantifiers of this form are monotone increasing and satisfy extension—once the quantifier is true in a model, adding new elements to A or $U - A$ will not change its logical value.⁵ The number triangle representation gives us always a downward triangle starting in a point $(0, m)$. How do those quantifiers behave in the Muddy Children situation?

Observation 1 *Let us take a Muddy Children situation, with n the number of children, $m \leq n$ the number of muddy children. The Muddy Children puzzle with the background assumption ‘At least k of you have mud on your forehead’ can be solved in $m - (k - 1)$ steps, where $k \leq m$.*

The proposition is verified in Figure 4. The number at the coordinates (c, m) says how many steps the *muddy* children need in order to converge to knowledge about their status (immediately after that all children know their status). The numbers of steps needed to solve the puzzle form a triangle, with the values increasing horizontally to the right. When increasing the parameter k in the quantifier ‘At least k ’ the whole triangle simply moves to the right and downwards.

Using a similar background assumption with inner negation At least k of you do not have mud on your foreheads also makes the puzzle solvable. ‘At least k not’ behaves as ‘At least k ’, but depends on the number of clean children. In general, inner negation works this way for other quantifiers.

A simplifying observation about a similar class of upward monotone quantifiers that satisfy extension is as follows:

Observation 2 *Let us take a Muddy Children situation, with n the number of children, $m \leq n$ the number of muddy children. The Muddy Children puzzle with the background assumption ‘More than k of you have mud on your forehead’ can be solved in $m - k$ steps, where $k \leq m$.*

Decreasing Quantifiers Let us now consider another natural class, downward monotone quantifiers that satisfy extension: ‘At most k ’, where $k \in \mathbb{N}$. In the puzzle this takes form of the announcement: At most k of you have mud on your foreheads.

Observation 3 *Let us take a Muddy Children scenario, with n the number of children, $m \leq n$ the number of muddy children. If $n > k$ then the Muddy Children puzzle with the background assumption ‘At most k of you have mud on your forehead’ can be solved in $(k + 1) - m$ steps. If $n \leq k$ the situation is not solvable.*

⁵In the case of (1,1) quantifiers this property corresponds to upward monotonicity in the left argument, which is also called *persistence* [7].

In Figure 5, like in the case of increasing quantifiers, we provided a triangle with the numbers of steps needed for solving the puzzle in respective cases. By doing this we indicate how the situation changes with the parameter k . We can observe that the numbers of steps needed to solve the puzzle form a block, with the values increasing horizontally to the left. When increasing the parameter k in the quantifier ‘At least k ’, the whole block moves to the right and downwards revealing the next column on the left. Also, in case when the parameter k in the quantifiers is larger or equal to the number of muddy children, the puzzle is not solvable. When the block of numbers moves downward together with k , it leaves a trace consisting of question marks that correspond to the unsolvable situations.

Cardinal and Parity Quantifiers Some kinds of quantifiers allow one-step immediate solvability for all agents. Taking into consideration what they already know, the announcement gives them full certainty about their state. This takes place for example when the number of muddy children is explicitly announced with the use of the quantifier ‘Exactly k ’, where $k \in \mathbb{N}$. The announcement of: Exactly k of you have mud on your foreheads always leads to immediate answers.

Observation 4 *Every Muddy Children scenario with a background assumption of the form ‘Exactly k ’ is solvable in 1 step.*

There are other, more interesting quantifiers with this property, e.g., divisibility quantifiers: ‘A number divisible by k ’, where $k \in \mathbb{N}$. An example of such an announcement for $k = 2$ is: An even number of you have mud on your foreheads. A relevant fact is as follows.

Observation 5 *Let us take a Muddy Children scenario. The Muddy Children puzzle with the background assumption ‘The number of you that have mud on your forehead is $\ell \bmod k$ ’, for any $\ell, k \in \mathbb{N}$, can be solved in 1 step.*

In Figure 6 the columns that include solvable scenarios are isolated and consists only of 1s. Moreover, if the number k in the quantifier ‘Divisible by k ’ increases the gaps between the columns.

Proportional Quantifiers Proportional quantifiers indicate the ratio between the number of elements in the predicate and the total number of elements. The first that comes to mind is ‘Exactly $\frac{1}{k}$ ’, where $k \in \mathbb{N}$. Update with this information will be survived by cardinalities that are divisible by k . In those situations, where $|A| = k \times \ell$, for some $\ell \in \mathbb{N}$, it is equivalent to the cardinal quantifier ‘Exactly ℓ ’. However, there are also more interesting cases of upward monotone proportional quantifiers. Such class is, e.g., ‘More than $\frac{1}{k}$ ’, where $k \in \mathbb{N}$. An example of such announcement could be: ‘Most of you have mud on your foreheads.’

If we agree to interpret ‘Most’ as ‘More than half’, then the solvability of the Muddy Children puzzle with this quantifier is depicted on the left in Figure




Fig. 7. ‘More than half’ and ‘More than one third’

7. The table on the right shows the pattern for the quantifier ‘More than one third’.

The patterns in Figure 7 might at first sight seem complex, but as a matter of fact it is quite easy to observe that the pattern consists of smaller parts resembling simple increasing quantifiers that satisfy extension (see Section 2.4). In fact these muddy situations are reducible to those given by quantifiers ‘More than k ’. In a given situation, when $|U| = n$, ‘More than $\frac{1}{k}$ ’ is of course equivalent to ‘More than $\frac{n}{k}$ ’.

Observation 6 *The epistemic quantifier ‘More than $\frac{1}{k}$ ’ consists of intervals q_0, q_1, \dots such that:*

- (1) q_0 consists of $k - 1$ rows in the triangle, and for $i > 0$, q_i consists of k rows.
- (2) q_i is the segment of size k of the table for ‘More than i ’ starting in the i -th row.

The number of steps needed to solve this puzzle is then characterized in the following way.

Observation 7 *Let us take a Muddy Children situation, with n the number of children, $m \leq n$ number of muddy children. The Muddy Children puzzle with the background assumption ‘More than $\frac{1}{k}$ of you have mud on your forehead’ can be solved in $\lceil m - \frac{n}{k} \rceil$ steps.*

3 Generalized Top Hat Puzzle

In the Muddy Children puzzle the agents have symmetrical observation powers. Each of them sees all agents except himself. The fact of the symmetry of the situation is common knowledge among the children and hence quite liberal background assumptions lead to the convergence to knowledge for every agent.

Things are different in the Top Hats puzzle. The observation is not symmetric; assuming there are n agents, the first one sees $n - 1$, the second $n - 2$, etc., and in the end there is always one agent that sees no one. Moreover in the solution of the puzzle it is required that they not only announce whether they know, but that they also explicitly say *what* they know about themselves.

Let us simplify the scenario. Assume that the first prisoner takes a role similar to the one of the father in Muddy Children puzzle. He simply announces the parity of the set of prisoners wearing black hats. Assume there are four agents, standing single file in the following order: a_1, a_2, a_3, a_4 , and that a_1 sees all other agents, a_2 sees a_3 and a_4 , etc. Moreover, let us assume the following hat distribution: a_1 and a_4 are wearing white hats and the others—black ones. Let us get back to our number triangle representation and let us interpret each pair (c, m) as standing for any situation in which there are c agents with white hats and m with black ones. Picture 8 indicates the actual world $(2,2)$, and observations of each agent (shaded states). Now, the truthful parity announcement is




Fig. 8. The Number-triangle Top-hat puzzle model

given (recall the number triangle representation of the quantifier from Fig. 6). The announcement divides the uncertainty range of each agent. However, only in the case of agent a_1 the announcement is fully informative in the first stage. He knows the parity of the overall number of black hats (by announcement) and he knows the parity of the number of black hats within the set of all agents except him. It is hence enough for him to conclude the color of his hat. When he announces it, the scenario repeats for agent a_2 , he now knows the parity of the overall number of black hats (by announcement) and he knows the parity of the number of black hats within the set of all agents except him; he can conclude the the color of his hat. Let us see why it was not possible in the first place. In the beginning, agent's a_2 observational state is $(1,1)$, he considers two immediate extensions of his states $(2,1)$ and $(1,2)$. However, each of the two extensions determine an observational state of agent a_1 and hence, from the perspective of agent a_2 , allow two possible observational states for a_1 .

3.1 Solvability

Which quantifiers are suitable for this kind of reasoning? Besides parity, if the situation allows, they could also convey all information needed using, for instance, ‘all’, ‘no’ or ‘all but k ’. Let us characterize the range of possibilities.

Definition 6. *We say that Q of type (1) has bounded thickness if there is a finite bound $k \in \mathbb{N}$ such that at any level in the number triangle Q allows not more than k consecutive $+$'s in the row.*

In particular, Q is 1-thick (has the thickness bounded by 1) iff the following conditions are satisfied:

- (1) If $(c, m) \in Q$ and $c > 0$ then $(c - 1, m + 1) \notin Q$, and
- (2) If $(c, m) \in Q$ and $m > 0$ then $(c + 1, m - 1) \notin Q$.

To prove the characterization theorem we need the following lemma:

Lemma 1. *Assume there are n agents: a_1, \dots, a_n (ordered according to decreasing observation power) and that the actual world is (c, m) such that $c + m = n$. Then at k -th stage of the puzzle agent a_k 's uncertainty range is either between $(c - 1, m + 1)$ and (c, m) or $(c + 1, m - 1)$ and (c, m) .*

Proof. First of all note that for any round k agent a_k 's uncertainty never includes more than 2 neighboring possibilities. It is so because he is only unsure about the color of his own hat. Moreover, observe that one of the possibilities have to be the actual world, if only agents are truthful and perfect reasoners.

Theorem 2. *A Top Hat situation with Q as the background assumption is solvable for any number of agents and any distribution of hats iff Q is 1-thick.*

Proof. From left to right. Assume that the puzzle with quantifier Q is solvable for any number of agents. Then take any (c, m) . By the structure of the puzzle $(c, m) \in Q$ (the announcement is truthful). Now, assume towards contradiction that (1) $(c - 1, m + 1) \in Q$ or (2) $(c + 1, m - 1) \in Q$. Hence, if (1) then the puzzle is not solvable for the distribution of hats in which a_1 has a white hat, and if (2) the puzzle is not solvable, analogously, for a_1 having a black hat.

From right to left. Assume that Q is 1-thick. Take any situation $(c, m) \in Q$, then $(c - 1, m + 1) \notin Q$ and $(c + 1, m - 1) \notin Q$. Therefore, by Lemma 1 for any round k of the puzzle agent a_k has uncertainty either between (1) $(c - 1, m + 1)$ and (c, m) or (2) $(c + 1, m - 1)$ and (c, m) . In both cases the worlds different than (c, m) are not in Q and hence the announcement eliminates them as a possibility.

The property of thickness, as far as we know, has not been defined in the literature. Through its name bounded thickness bears resemblance to the notion of finite thickness used in the domain of inductive inference [9]. Finite thickness is a sufficient condition for identifiability in the limit [10]. In our terms (consistent with [11]) a class of quantifiers \mathcal{Q} has finite thickness if for every finite model M there are only finitely many quantifiers in \mathcal{Q} that include M . Bounded thickness

does not imply finite thickness. Note however, that this is not enough to conclude that quantifiers satisfying bounded thickness are not learnable.

Bounded thickness is related to the oscillation properties of generalized quantifiers [12]. Q of type (1) has bounded oscillation property if there is a finite bound $m \in \mathbb{N}$ such that at any level in the number triangle Q switches from $+$ to $-$, or vice versa, not more than m times. For example, any monotone quantifier has bounded oscillation with $m=1$ and the quantifier ‘either between 3 and 5 or more than 8’ has bounded oscillation with $m = 3$ (see [7], Ch. 14.4). Obviously, the quantifier does not have to satisfy bounded oscillation to make the puzzle solvable (in the original formulation it includes unboundedly oscillating parity quantifier). A different kind of restriction is required, namely, that the ‘wideness’ of each segment of states belonging to the quantifier relation within one level is equal to one, i.e., if a model belongs to the quantifier none of its neighboring models belong to it. We call that property 1-thickness. It has to do with the binary structure of the situation—the quantifier must eliminate one of two states for every uncertainty range.

Various invariance properties of quantifiers are strongly present in the domain of generalized quantifier theory, especially because of their implications for linguistic universals of natural language determiners [13]. We leave it as an open question what is the exact relation of those issues to thickness, it seems however that this property gives a quantifier a ‘high resolution’ in distinguishing similar situation and, hence, implies a significant informational content. In order to relate to the Muddy Children puzzle let us note here that 1-thickness is a stronger condition than positive variation:

Observation 8 *For any Q , if it is 1-thick then it satisfies VAR^+ .*

Moreover, 1-thickness of the Top Hat puzzle quantifiers justifies the counter-intuitiveness of the whole scenario. Recall, that according to a popular view [13] simple natural language determiners are expected to satisfy continuity (a property that characterizes exactly conjunctions of an increasing and a decreasing quantifier, see e.g. Ch. 5.3 in [7]). However, many non-trivial continuous quantifiers are not 1-thick.

Definition 7. *A type (1) quantifier Q_M is continuous (CONT) iff for any M and all $A', A'' \subseteq M$, if $Q_M(A'), Q_M(A'')$, and $A' \subseteq A \subseteq A''$, then $Q_M(A)$.*

Observation 9 *Assume (for non-triviality) that $Q_M(A'), Q_M(A'')$ and that Q_M is CONT then Q_M cannot be 1-thick.*

Therefore, in a sense quantifiers satisfying bounded thickness are strongly non-monotone.

4 Conclusions and Outlook

In this paper we analyzed and generalized the popular Muddy Children puzzle. We investigated the epistemic and informational properties of the situation under

various quantifier announcements. We characterized the solvability of the puzzle as dependent on invariance properties of the quantifier involved in the public announcement. In particular, we have shown that a weaker version of variance, VAR^+ , is sufficient to trigger a successful epistemic reasoning. We applied a similar analysis, based on the number triangle representation of quantifiers, to the Top Hat puzzle in order to explicitly account for the difference in the informational power of various quantifiers. Here, it is not enough that a quantifier is active to solve every instance of the puzzle. The announcement needs to make sure that any two models neighboring each other in the number triangle do not belong to the quantifier simultaneously.

We also proposed a concise modeling of epistemic scenarios that may be attractive in all those applications where an agent’s internal representation of the problem is crucial, like cognitive science or designing multi-agent systems in the domain of artificial intelligence. One of the main aims of applying logic in those disciplines is to model possible inferential strategies of an agent. An immediate plausibility test is the complexity of the proposed representation. The classical epistemic logic based modeling of such scenarios usually assume an extensive representation including all possibilities (see e.g. [2]). In our case, the size of the models is clearly connected to the properties of generalized quantifiers—their isomorphism closure can increase the informational power of a message relatively to the observational powers of the agents. Such informational ‘shortcuts’ are rarely taken into account in the epistemic literature.

Our study provides an additional interesting link with formal epistemology. An agent in the Muddy Children puzzle can be seen as a scientist who tries to inductively decide a hypothesis, tries to discover what the actual world is like (see e.g. [14], [15]). Our analysis shows that even if the agents have limited observational capacities, the presence and interconnection with other scientists doing similar research can influence the discovery in a positive way (cf. [16]). In this sense the following paper can be positioned among various attempts to study the *learnability* issues related to generalized quantifiers. One perspective here is to study formal models of quantifier acquisition, for example in Gold paradigm [10] it is interesting to ask whether certain classes of quantifiers are identifiable in the limit (see [11]). Similarly, one can be interested in more algorithmic approach and asks about the complexity of procedures learning certain classes of quantifiers over finite models [17]. Another angle is to study the computational complexity of deciding whether a given finite model satisfy various classes of natural language quantifiers [18]. Those two perspectives of identifying and verifying a quantifier hypothesis can be combined by studying the problems on an inductively given information about the actual world [19]. Finally, in this work we investigate how the nature of the world may be identified by a group of agents on the basis of some quantifier information and restricted communication. There are many directions of follow-up research. One attractive idea is to develop a general framework for investigating the informational properties of generalized quantifiers in single- and multi-agent learning contexts that would account for a wide variety of scenarios.

References

1. Littlewood: A mathematician's miscellany. London: Meuthen (1953)
2. Van Ditmarsch, H., Van der Hoek, W., Kooi, B.: Dynamic Epistemic Logic. Springer Netherlands (2007)
3. Gierasimczuk, N., Szymanik, J.: A note on a generalization of the muddy children puzzle. In: Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2011), Groningen, The Netherlands, July 12–14, 2011, ACM (2011)
4. Gierasimczuk, N.: Knowing One's Limits. Logical analysis of inductive inference. PhD thesis, Universiteit van Amsterdam (2010)
5. Mostowski, A.: On a generalization of quantifiers. *Fundamenta Mathematicae* **44** (1957) 12–36
6. Van Benthem, J.: *Essays in Logical Semantics*. D. Reidel, Dordrecht (1986)
7. Peters, S., Westerståhl, D.: *Quantifiers in Language and Logic*. Oxford University Press, Oxford (2006)
8. Van Benthem, J.: Questions about quantifiers. *Journal of Symbolic Logic* **49**(2) (1984) 443–466
9. Angluin, D.: Finding patterns common to a set of strings. *Journal of Computer and System Sciences* **21**(1) (1980) 46–62
10. Gold, E.: Language identification in the limit. *Information and Control* **10** (1967) 447–474
11. Tiede, H.J.: Identifiability in the limit of context-free generalized quantifiers. *Journal of Language and Computation* **1** (1999) 93–102
12. Väänänen, J., Westerståhl, D.: On the Expressive Power of Monotone Natural Language Quantifiers over Finite Models. *Journal of Philosophical Logic* **31**(4) (2002) 327–358
13. Barwise, J., Cooper, R.: Generalized quantifiers and natural language. *Linguistics and Philosophy* **4** (1981) 159–219
14. Jain, S., Osherson, D., Royer, J.S., Sharma, A.: *Systems that Learn*. MIT Press, Chicago (1999)
15. Kelly, K.: *The Logic of Reliable Inquiry*. Oxford University Press, Oxford (1996)
16. Dégremont, C., Gierasimczuk, N.: Finite identification from the viewpoint of epistemic update. *Information and Computation* **209** (2011) 383–396
17. Gierasimczuk, N.: The problem of learning the semantics of quantifiers. In: *Lecture Notes in Artificial Intelligence*. Volume 4363., Springer (2007) 117–126
18. Szymanik, J.: Computational complexity of polyadic lifts of generalized quantifiers in natural language. *Linguistics and Philosophy* **33** (2010) 215–250
19. Gierasimczuk, N.: Identification through inductive verification. In: *Lecture Notes in Computer Science*. Volume 5422., Springer-Verlag (2009) 193–205