

Anonymous Public Announcements

Thomas Ågotnes

University of Bergen, Bergen, Norway

Shanxi University, Taiyuan, China

thomas.agotnes@uib.no

Rustam Galimullin

University of Bergen, Bergen, Norway

Rustam.Galimullin@uib.no

Ken Satoh

Center for Juris-Informatics, Tokyo, Japan

ksatoh@nii.ac.jp

Satoshi Tojo

Asia University, Tokyo, Japan

tojo.satoshi@gmail.com

Abstract

We formalise the notion of an *anonymous public announcement* in the tradition of public announcement logic. Such announcements can be seen as in-between a public announcement from “the outside” (an announcement of ϕ) and a public announcement by one of the agents (an announcement of $K_a\phi$): we get more information than just ϕ , but not (necessarily) about exactly who made it. Even if such an announcement is *prima facie* anonymous, depending on the background knowledge of the agents it might reveal the identity of the announcer: if I post something on a message board, the information might reveal who I am even if I don’t sign my name. Furthermore, like in the Russian Cards puzzle, if we assume that the announcer’s *intention* was to stay anonymous, that in fact might reveal more information. In this paper we first look at the case when no assumption about intentions are made, in which case the logic with an anonymous public announcement operator is reducible to epistemic logic. We then look at the case when we assume common knowledge of the intention to stay anonymous, which is both more complex and more interesting: in several ways it boils down to the notion of a “safe” announcement (again, similarly to Russian Cards). Main results include formal expressivity results and axiomatic completeness for key logical languages.

1 Introduction

Taken at the face value, the title of this paper seems to be an oxymoron. Indeed, if “public announcement” is taken literally, as in an agent *saying* something in front of everyone else, it will not be anonymous. However, anonymous public communication is almost ubiquitous in our day-to-day lives. Think of posts on social media and message boards done under a username instead of a real name of the poster. Or an anonymous letter to an editor of a news outlet. Other examples include anonymous emails, transactions on a blockchain, whistle blower reports, and even cultural artifacts created anonymously under an alias, like Elena Ferrante and MF DOOM.

The type of anonymity we are interested here focuses on *action anonymity*, i.e. the inability of an attacker to identify who performed a given action (also sometimes referred to as *unlinkability* in the literature [31]). This is in contrast to *data anonymity*, i.e. the inability of an attacker to know the identity of a subject in an anonymised database, e.g. in medical records (see, e.g., [16]). One of the standard requirements to both types of anonymity is that they satisfy *k*-anonymity [34], which intuitively means that a data record or an action cannot be distinguished from at least $k - 1$ other records or actions. It is clear that in the case of public communication, we should have at least 3-anonymity. Indeed, if a public announcement is so specific that it could be done only by two agents (and these two agents know it), then the non-announcing agent would be able to deduce the identity of the announcer.

In this paper we formalise anonymous public announcements inspired by *public announcement logic* (PAL) [32], an extension of multi-agent *epistemic logic* with constructs of the form $[\phi!]\psi$ intuitively meaning that after ϕ is truthfully announced, ψ is true. In PAL the announced formula ϕ does not have to actually be known by any agent in the system – the identity of the announcer is left out of the picture. If the announcer indeed is one of the agents a in the system, the announcement in fact contains *more* information: in that case it would be modelled by the announcement $K_a\phi!$. In this paper we formalise *anonymous* public announcements, conceptually somewhere in-between $\phi!$ and $K_a\phi!$ – we get *more* information than just ϕ but *less* than $K_a\phi$ for a specific agent a . We thus first introduce modalities of the form $[\psi\dagger]$ where $[\psi\dagger]\phi$ that are read as “after ψ is *pseudo-anonymously* announced, ϕ is true no matter who the announcer was”. Such announcements are pseudo-anonymous because we do not assume that the announcer necessarily *intends* to stay anonymous. Indeed, the reader may have noticed that not all the examples of anonymous public communication we mentioned above are guaranteed to preserve the identity of the announcer. For example, if the announced information is known only to two agents (and the non-announcing agent knows it), then the identity of the announcer will not be a secret to the non-announcing agent. In the literature, such a scenario is called “background knowledge attack” [25]. Thus, we also formalise *intentional* (or *safe*) *anonymous announcements* by introducing constructs $[\psi\dagger]\phi$ meaning “after the safe announcement of ψ , ϕ is true no matter who the announcer was”. In order to capture safety, we employ the idea of (at least) 3-anonymity, and

agents intending to stay anonymous should know that their announcement is safe. A bit more formally, we introduce the safety modality $\blacktriangle\psi$ which, similarly to common knowledge, is defined via a fixpoint. Hence, $\blacktriangle\psi$ means that “ ψ is true, and that there is a group of three agents that know that ψ , and there is a group of three agents that know that there is a group of three agents that know that ψ and so on”. In other words, safe anonymous announcements guarantee the anonymity of announcers.

Reasoning about anonymity based on (variants of) epistemic logic has been studied in [35, 19], with k -anonymity being discussed in [19]. Building on the runs-and-systems approach of [19], further extensions focusing on privacy and onymity [36, 37] and electronic voting [26]. A knowledge-based approach to data anonymity was presented in [23]. Other logical approach to anonymity and privacy include [33, 22, 40, 11, 30]. The themes of anonymity and privacy are also related to the research on secrets in multi-agent systems (see, e.g., [20, 29, 28, 39]). None of these approaches model anonymous *announcements*. Finally, an approach to anonymity based on *dynamic epistemic logic* (DEL) [14] was presented in [38], where the authors consider scenarios of secret communication between agents in a system. Such secret communication is captured by private announcements, as special type of *action models* [10]. In this setting, verifying whether secret communication remained secret boils down to model checking an epistemic formula in the resulting updated model, i.e. the model that is obtained after an application of a dynamic operator.

While our approach is also DEL-based, in our work we focus on anonymous *public* communication, as opposed to *private* communication. Moreover, we tackle the issue of intentional anonymity, where the announcing agent is guaranteed to stay anonymous. To this end, we introduce a novel safety modality that, to the best of our knowledge, has never been considered in the literature before.

We start out by introducing the machinery of EL and DEL in Section 2. In Section 3, we study pseudo-anonymous announcements, and in particular show that while the resulting logic is as expressive as EL, standard and pseudo-anonymous announcements are update incomparable. After that, in Section 4, we formalise intentional anonymous announcements using the safety modality, as well as show that, perhaps surprisingly, safety is all one needs: logics with safe announcements are as expressive as EL extended with the safety modality. Moreover, we provide a sound and complete axiomatisation of the latter. Finally, in Section 5 we discuss further research directions.

2 Background

Let N be a finite set of *agents*, and P be a countable set of *propositional variables*.

All logics that we are dealing with in this paper are interpreted on epistemic models.

Definition 1 An (epistemic) model is a triple $M = (S, \sim, V)$, where S is a non-empty set of states, $\sim: N \rightarrow 2^{S \times S}$ is an equivalence relation for each $i \in N$, and $V: P \rightarrow 2^S$ is the valuation function. For $s \in S$, a pair M, s is called a pointed (epistemic) model.

Before we define some basic logical languages, we need to introduce a special kind of models that will be used in syntax.

Definition 2 Let \mathcal{L} be a language defined over the signature $\langle N, P \rangle$. An action model is a triple $\mathbf{M} = (\mathbf{S}, \sim, \text{pre})$, where \mathbf{S} is a non-empty set of states, $\sim: N \rightarrow 2^{S \times S}$ is an equivalent relation for each $i \in N$, and $\text{pre}: \mathbf{S} \rightarrow \mathcal{L}$ is the precondition function. For $\mathbf{s} \in \mathbf{S}$, we will call a pair \mathbf{M}, \mathbf{s} a pointed action model.

Definition 3 Languages of epistemic logic (\mathcal{L}_K), public announcement logic ($\mathcal{L}_!$), and action model logic (\mathcal{L}_\otimes) are defined by the following BNFs:

$$\begin{aligned} \mathcal{L}_K \ni \phi &::= p \mid \neg\phi \mid (\phi \wedge \psi) \mid K_i\phi \\ \mathcal{L}_! \ni \phi &::= p \mid \neg\phi \mid (\phi \wedge \psi) \mid K_i\phi \mid [\phi]\phi \\ \mathcal{L}_\otimes \ni \phi &::= p \mid \neg\phi \mid (\phi \wedge \psi) \mid K_i\phi \mid [\pi]\phi \\ \pi &::= (\mathbf{M}, \mathbf{s}) \mid \pi \cup \pi \end{aligned}$$

where $p \in P$, $i \in N$, and (\mathbf{M}, \mathbf{s}) is a pointed action model with a finite set of states \mathbf{S} , and such that for all $\mathbf{s} \in \mathbf{S}$, precondition $\text{pre}(\mathbf{s})$ is some $\phi \in \mathcal{L}_\otimes$ that was constructed in a previous stage of the inductively defined hierarchy.

Having defined models and languages, we are now ready to provide the definition of the semantics of the aforementioned logics.

Definition 4 Let $M, s = (S, \sim, V)$ be a model, $p \in P$, $i \in N$, and (\mathbf{M}, \mathbf{s}) and be an action model.

$$\begin{aligned} M, s \models p & \quad \text{iff} \quad s \in V(p) \\ M, s \models \neg\phi & \quad \text{iff} \quad M, s \not\models \phi \\ M, s \models \phi \wedge \psi & \quad \text{iff} \quad M, s \models \phi \text{ and } M, s \models \psi \\ M, s \models K_i\phi & \quad \text{iff} \quad M, t \models \phi \text{ for all } t \in S \text{ such that } s \sim_i t \\ M, s \models [\psi!]\phi & \quad \text{iff} \quad M, s \models \psi \text{ implies } M, s^{\psi!} \models \phi \\ M, s \models [\mathbf{M}, \mathbf{s}]\phi & \quad \text{iff} \quad M, s \models \text{pre}(\mathbf{s}) \text{ implies } (M \otimes \mathbf{M}, (s, \mathbf{s})) \models \phi \\ M, s \models [\pi \cup \rho]\phi & \quad \text{iff} \quad M, s \models [\pi]\phi \text{ and } M, s \models [\rho]\phi \end{aligned}$$

Given M , we will write $[[\phi]]_M$ for the set $\{s \in S \mid M, s \models \phi\}$.

The updated model $M^{\phi!}$ is $(S^{\phi!}, \sim^{\phi!}, V^{\phi!})$, where $S^{\phi!} = [[\phi]]_M$, $\sim_i^{\phi!} = \sim_i \cap (S^{\phi!} \times S^{\phi!})$ for all $i \in N$, and $V^{\phi!}(p) = V(p) \cap [[\phi]]_M$ for all $p \in P$.

Updated model $M \otimes \mathbf{M}$ is (S', \sim', V') , where $S' = \{(s, \mathbf{s}) \mid s \in S, \mathbf{s} \in \mathbf{S}, M, s \models \text{pre}(\mathbf{s})\}$, $(s, \mathbf{s})R'_i(t, \mathbf{t})$ iff $s \sim_i t$ and $\mathbf{s} \sim_i \mathbf{t}$, and $(s, \mathbf{s}) \in V'(p)$ iff $s \in V(p)$.

We call a formula ϕ valid, or a validity, if for all M, s it holds that $M, s \models \phi$.

Definition 5 Let $M^1 = (S^1, \sim^1, V^1)$ and $M^2 = (S^2, \sim^2, V^2)$ be two epistemic models. We say that M^1 and M^2 are bisimilar (denoted $M^1 \rightleftharpoons M^2$) if there is a non-empty relation $Z \subseteq S^1 \times S^2$, called bisimulation, such that for all sZt , the following conditions are satisfied:

Atoms for all $p \in P$: $s \in V^1(p)$ if and only if $t \in V^2(p)$,

Forth for all $i \in N$ and $u \in S^1$ s.t. $s \sim_i^1 u$, there is a $v \in S^2$ s.t. $t \sim_i^2 v$ and uZv ,

Back for all $i \in N$ and $v \in S^2$ s.t. $t \sim_i^2 v$, there is a $u \in S^1$ s.t. $s \sim_i^1 u$ and uZv .

We say that M^1, s and M^2, t are bisimilar and denote this by $M^1, s \rightleftharpoons M^2, t$ if there is a bisimulation linking states s and t .

It is a standard result that $M^1, s \rightleftharpoons M^2, t$ implies $M^1, s \models \phi$ if and only if $M^2, t \models \phi$ for $\phi \in \mathcal{L}_K$ (see, e.g., [18]). In other words, epistemic logic is bisimulation invariant over epistemic models. We also have that $\mathcal{L}_!$ and \mathcal{L}_\otimes are bisimulation invariant over epistemic models [14, Chapter 5].

Definition 6 Let \mathcal{L}_1 and \mathcal{L}_2 be two languages, and let $\phi \in \mathcal{L}_1$ and $\psi \in \mathcal{L}_2$. We say that ϕ and ψ are equivalent, when for all models M_s : $M_s \models \phi$ if and only if $M_s \models \psi$.

If for every $\phi \in \mathcal{L}_1$ there is an equivalent $\psi \in \mathcal{L}_2$, we write $\mathcal{L}_1 \preceq \mathcal{L}_2$ and say that \mathcal{L}_2 is at least as expressive as \mathcal{L}_1 . We write $\mathcal{L}_1 \prec \mathcal{L}_2$ iff $\mathcal{L}_1 \preceq \mathcal{L}_2$ and $\mathcal{L}_2 \not\preceq \mathcal{L}_1$, and we say that \mathcal{L}_2 is strictly more expressive than \mathcal{L}_1 . Finally, if $\mathcal{L}_1 \preceq \mathcal{L}_2$ and $\mathcal{L}_2 \preceq \mathcal{L}_1$, we say that \mathcal{L}_1 and \mathcal{L}_2 are equally expressive and write $\mathcal{L}_1 \approx \mathcal{L}_2$.

The classic result in DEL is that $\mathcal{L}_! \approx \mathcal{L}_\otimes \approx \mathcal{L}_K$ [14, Chapter 8], i.e. that any formula with public announcements or action models can be reduced to an equivalent formula of epistemic logic.

While $\mathcal{L}_!$ and \mathcal{L}_\otimes are equally expressive, in some sense action models are more powerful than public announcements. Indeed, a public announcement of ψ can be modelled by the single-state action model $M^\psi = (\{s, \{s \sim_i s \mid i \in N\}, \text{pre}(s) = \psi\})$. At the same time, not every action model can be modelled by a public announcement: action models may result in the bigger updated models, while public announcements result in the updated models of the size of at most the size of the original one. We can capture this distinction formally.

Definition 7 Let \mathcal{L}_1 and \mathcal{L}_2 be two languages with dynamic operators. If for every update $[\alpha]$ of \mathcal{L}_1 there is an update $[\beta]$ of \mathcal{L}_2 such that for all M, s , the update of M, s with α is bisimilar to the update of M, s with β , then we say that \mathcal{L}_2 is at least as update expressive as \mathcal{L}_1 and write $\mathcal{L}_2 \preceq^U \mathcal{L}_1$. We write $\mathcal{L}_1 \prec^U \mathcal{L}_2$ iff $\mathcal{L}_1 \preceq^U \mathcal{L}_2$ and $\mathcal{L}_2 \not\preceq^U \mathcal{L}_1$, and we say that \mathcal{L}_1 is strictly less update expressive than \mathcal{L}_2 . If $\mathcal{L}_1 \preceq^U \mathcal{L}_2$ and $\mathcal{L}_2 \preceq^U \mathcal{L}_1$, we say that \mathcal{L}_1 and \mathcal{L}_2 are equally update expressive and write $\mathcal{L}_1 \approx^U \mathcal{L}_2$. Finally, if $\mathcal{L}_1 \not\preceq^U \mathcal{L}_2$ and

$\mathcal{L}_2 \not\leq^U \mathcal{L}_1$, we say that \mathcal{L}_1 and \mathcal{L}_2 are incomparable in update expressivity and write $\mathcal{L}_1 \not\leq^U \mathcal{L}_2$.

See [15] for a more thorough discussion on update expressivity as well as an overview of update expressivity results for some DELs, including $\mathcal{L}_! \prec^U \mathcal{L}_\otimes$.

3 Unintentional Anonymity

We are interested in anonymous public communication that we encounter in anonymous P2P networks, Internet forums, some blockchains, and so on. We can also think about the event of “all the agents come into the classroom, and see that someone (one of the agents) has written ϕ on the blackboard”. This announcement is fundamentally anonymous in the sense that the announcers’ identity is not explicitly revealed, but it might be *implicitly* revealed in certain situations depending on the agents’ background knowledge. A trivial example is if you write something on the blackboard that I know that only you know. We thus call such events *pseudo-anonymous announcements*.

Unlike a standard public announcement of ϕ by a , a pseudo-anonymous announcement of ϕ by a is no longer modeled by a public announcement of the formula $K_a\phi$, since the announcement is not signed by a . On the other hand, it is also not modeled by a public announcement of ϕ , since the event also contains information that *some* agent knows ϕ . The example “someone has written ϕ on the blackboard” hints to a public announcement of “someone knows ϕ ”, i.e., $\bigvee_{i \in N} K_i\phi$.

However, updating the model by removing states where $\bigvee_{i \in N} K_i\phi$ is false also does not work, because, e.g., it does not take into account the fact that the announcer knows who the announcer is. Indeed, this also shows that pseudo-anonymous announcements *cannot be modeled by public announcements at all, because the former are not deterministic*¹: it might be that a announced ϕ , and it might be that b announced ϕ , and these are two different model updates that need to be considered separately.

We will do exactly that, by first defining the updated model and then defining the semantics a new dynamic modality $[\phi^\dagger]\psi$ in terms of that. Similar to the standard public announcement operator $[\phi!]$, $[\phi^\dagger]\psi$ is intended to mean that ψ is true after ϕ is pseudo-anonymously announced. There is a very direct relationship to action models, which we make explicit in Section 3.1 below.

Definition 8 *The update of epistemic model $M = (S, \sim, V)$ by the pseudo-anonymous announcement of ϕ is the epistemic model $M^{\phi^\dagger} = (S', \sim', V')$ where:*

- $S' = \{(s, a) \mid s \in S, a \in N \text{ and } M, s \models K_a\phi\}$
- $(s, a) \sim'_c (t, b)$ iff $s \sim_c t$ and $a = c$ iff $b = c$
- $V'(s, a) = V(s)$

¹We will make this observation more precise in Proposition 2.

Intuitively, in the updated model it is common knowledge that someone had (truthfully) announced/posted/written on the blackboard that ϕ . A state (s, a) in the updated model corresponds to that someone being a , in state s of the original model. The precondition is that a knows ϕ , corresponding to the assumption that anonymous public announcements are truthful. An agent can only discern between a situation (s, a) where a made the pseudo-anonymous announcement and a situation (t, b) where a different agent b did, if she could already discern between s and t before the announcement or she is neither a nor b .

Definition 9 *The language of pseudo-anonymous public announcement logic \mathcal{L}_\dagger is recursively defined by the BNF*

$$\phi ::= p \mid \neg\phi \mid (\phi \wedge \phi) \mid K_i\phi \mid [\phi^\dagger]\phi$$

Intuitively, $[\phi^\dagger]$ is the event “ ϕ is pseudo-anonymously announced”. $[\phi^\dagger]\psi$ means that ψ is necessarily true after that event, no matter who the announcer was. The dual is defined as $\langle\phi^\dagger\rangle = \neg[\phi^\dagger]\neg\psi$, and intuitively means that there is an agent who can pseudo-anonymously announce ϕ such that ψ will be true afterwards.

$$M, s \models [\phi^\dagger]\psi \quad \text{iff} \quad \forall a \in N : M, s \models K_a\phi \text{ implies } M^{\phi^\dagger}, (s, a) \models \psi$$

For an example, see Figure 1. Let M be the epistemic model in the middle. The figure also show three different updates related to announcements about p : someone outside the system announced p (top); it was announced that someone inside the system knows p (bottom right); someone inside the system pseudo-anonymously announced that they know p (bottom left).

Note that in this example:

- All the three updates are different.
- $M, s \models [p^\dagger]\psi$ holds iff ψ holds in both states (s, a) and (s, c) in the updated model to the bottom left in the figure.
- Thus we have for example that $M, s \models [p^\dagger]\hat{K}_b(r \wedge K_a\neg q)$: after the announcement b considers it possible that the initial state was s and that a has learned that q is false. This holds because b considers it possible it was c who announced p (state (s, c) in the updated model), in which case a has learned that we were not in u . Note that this formula does not hold after the two other updates: $M, s \not\models [p!]\hat{K}_b(r \wedge K_a\neg q)$ and $M, s \not\models [\bigvee_{i \in N} K_i p!]\hat{K}_b(r \wedge K_a\neg q)$.
- The difference between the updates ϕ^\dagger and $\bigvee_{i \in N} K_i\phi!$ is related to the familiar “de re”/“de dicto” distinction: in the latter case we update with “in every state there exists someone who knows ϕ ”, in the former with “there exists someone who knows ϕ in every state”.

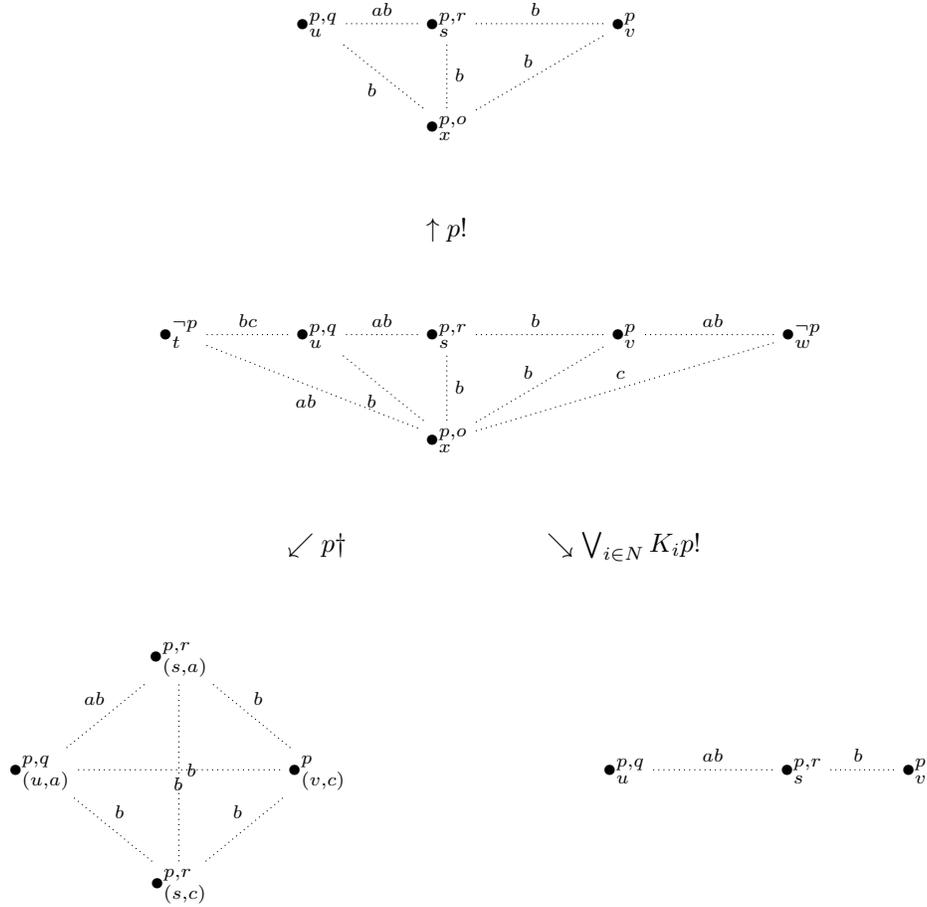


Figure 1: Three-agent epistemic model (middle) and its update after the event “somebody pseudo-anonymously said p ” (bottom left) as well as the update after the public announcement of “somebody knows ϕ ” (bottom right) and the update after the public announcement “ ϕ is true” (top). (Assume transitive closure, not all edges are shown.)

- We also have that $M, s \not\models [p^\dagger]K_a r$: if it was a herself who made the pseudo-anonymous announcement (state (s, a) in the updated model), she wouldn't learn that we were in s .

Pseudo-anonymous announcements are not necessarily *truly* anonymous, in all situations w.r.t. all agents, because depending on the background information of the agents they might reveal the identity of the announcer. Continuing with the example:

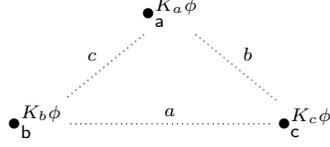


Figure 2: The three-agent pseudo-anonymous action model.

- a 's pseudo-anonymous announcement of p is truly anonymous w.r.t. b , in the epistemic state M, s . Indeed, b considers it possible that it was c who announced p . We see this explicitly in the updated model: b considers two states possible where a announced p and two where c did it.
- a ' pseudo-anonymous announcement of p is not truly anonymous w.r.t. c , in the epistemic state M, s . Intuitively, this is because c knows that the only other agent besides herself who can announce p is a . Formally, this can be seen from the updated model: in the state (s, a) , c knows that it was a who made the announcement.

3.1 Relationship to Action Models

As mentioned above, there is a close relationship with action model logic. In fact, the semantics of ϕ^\dagger given above is equivalent to a certain class of action models, *pseudo-anonymous action models*, as we now show.

Definition 10 *The pseudo-anonymous action model for N agents and formula ϕ is the action model $M_\phi^N = (S, \sim, \text{pre})$ where*

- $S = N$
- $a \sim_b c$ iff $(a = c \Leftrightarrow b = c)$
- $\text{pre}(a) = K_a \phi$

Intuitively, the pointed pseudo-anonymous action model (M_ϕ^N, a) corresponds to a writing ϕ on the blackboard. The precondition is that a knows ϕ . Agent c cannot discern between events a and b if and only if either $a = b$, or $a \neq c$ and $b \neq c$. The three-agent pseudo-anonymous event is illustrated in Figure 2.

A pseudo-anonymous announcement now corresponds to the *union* of events (M_ϕ^N, a) for all agents a . The proof of the following is straightforward by the definition of semantics.

Proposition 1 *For any pointed epistemic model M, s and formulas ϕ and ψ , $M, s \models [\phi^\dagger]\psi$ iff $M, s \models \left[\bigcup_{i \in N} (M_\phi^N, i) \right] \psi$.*

From this we get that \mathcal{L}_\dagger is no more expressive than \mathcal{L}_K via the known fact that $\mathcal{L}_\otimes \approx \mathcal{L}_K$.

3.2 Update expressivity

Given that the logic of pseudo-anonymous announcements is as expressive as epistemic logic, we show that pseudo-anonymous announcements are still unique dynamic operators. To this end, we claim that $\mathcal{L}_\dagger \not\approx^U \mathcal{L}_!$ and $\mathcal{L}_\dagger \not\approx^U \mathcal{L}_\otimes$.

Proposition 2 \mathcal{L}_\dagger and $\mathcal{L}_!$ are update incomparable.

Proof: Consider model M in Figure 3 and public announcement $p!$. The result of the public announcement of $p!$ in M is the model $M^{p!}$ consisting of the single state s satisfying p .

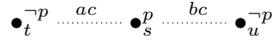


Figure 3: Epistemic model M .

To see that there is no pseudo-anonymous announcement ϕ^\dagger that will result in a model bisimilar to $M^{p!}$, it is enough to notice that none of the agents in model M has the ability to remove $\neg p$ -states due to precondition $K_i\phi$ for pseudo-anonymous announcements.

To show that there are pseudo-anonymous announcements that cannot be captured by public announcements, we turn back to the model M in the centre of Figure 1 and its update M^{p^\dagger} depicted at bottom left of the same figure. Moreover, consider formula $\phi := r \wedge \widehat{K}_a q \wedge \widehat{K}_b K_a(\neg q \wedge r) \wedge K_b p$. This verbose formula is true in $M^{p^\dagger}, (s, a)$, and intuitively states that the current state satisfies r (i.e. we are in an s -state), all reachable states satisfy p (via conjunct $K_b p$ and the fact that b 's relation is universal), that there is an a -reachable q -state (conjunct $\widehat{K}_a q$), and that there is a state, where a knows $\neg q \wedge r$ (state (s, c)).

Now we will argue that no public announcement made in M can result in an update satisfying ϕ . First, since we have $K_b p$ as a conjunct in ϕ , any update should preserve only p -states, i.e. a non-empty subset of $\{s, u, v, x\}$. State s should be preserved as it is the only one satisfying r . Similarly, state u should be in the updated model to satisfy $\widehat{K}_a q$. Now, the third conjunct of ϕ forces the existence of a state, where a knows $\neg q \wedge r$. The only state that satisfies $\neg q \wedge r$ is s . However, since we have to preserve u , it would not hold in s that $K_a(\neg q \wedge r)$. It is easy to check that none of the remaining options of adding x or v will lead to the satisfaction of $K_a(\neg q \wedge r)$ in s . Thus, there is no public announcement that can be made in M such that the updated model exists and satisfies ϕ . This implies that there is no public announcement update of M that is bisimilar to $M^{p^\dagger}, (s, a)$. \square

Corollary 1 \mathcal{L}_\dagger is strictly less update expressive than \mathcal{L}_\otimes .

This follows from the facts that (1) each anonymous announcement can be modelled by the corresponding action model (Proposition 1) and hence $\mathcal{L}_\dagger \preceq^U$

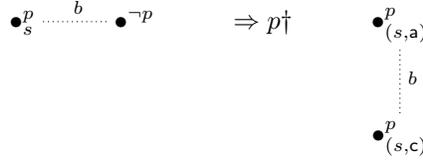


Figure 4: Three-agent epistemic model (top) and its update after a pseudo-anonymous announcement of p (bottom). Accessibility for agents a and c is the identity relation.

\mathcal{L}_\otimes ; (2) each public announcement can be modelled by a single-state action model ($\mathcal{L}_! \not\approx^U \mathcal{L}_\dagger$ from Proposition 2 and hence $\mathcal{L}_\otimes \not\approx^U \mathcal{L}_\dagger$).

4 Intentional Anonymity

We now turn to *intentional* anonymous announcements, where we assume that it is common knowledge that the announcement was intended to be anonymous. We want to capture that idea by strengthening the pre-condition on announcements, so that only “safe” announcements are allowed. We will discuss in detail what “safe” means, and it turns out to be quite subtle, but the intuition is that it is safe for an agent to make an announcement if somebody else could have made it. One of the subtleties is that the potential announcement by that “somebody else” also should be safe, so the definition has a recursive (or even circular!) flavour.

As a first simple example, consider the three-agent model in Figure 4. We have that both $M, s \models K_a p$ and $M, s \models K_c p$ so both a and c can announce p , and we see in the updated model that in either case b doesn’t know who did it. However, since there are only two agents who could make the announcement, “the other” agent learns who made it – if a makes the announcement, c knows that a did it, and vice versa. Thus, *2-anonymity is not enough*. This is because we modeling anonymity *within the system*, and not from the outside. So, we need at least *3-anonymity*: an announcement is only safe if at least three agents can make it. In the model in the figure, no safe announcements of p can be made.

Let’s move on to the model M in Figure 5, where also the updated model after the pseudo-anonymous announcement of p is shown. In state s , all of a , b and c know p and thus any one of them could have written it on the blackboard. The updated model reflects that no-one except the announcer knows who did it. However, in state s :

- a considers it possible that we are in state t , where it would not be safe for her to announce p since then b would know that it was her.
- Thus it is not safe for a to announce p in s .

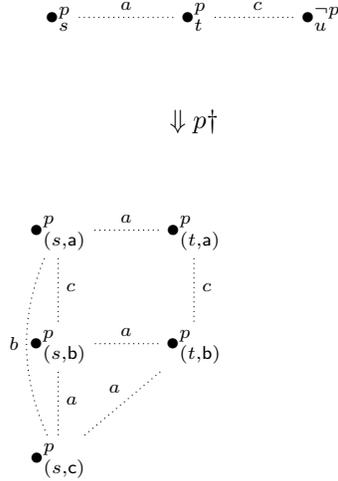


Figure 5: Three-agent epistemic model, with update. Accessibility for agent b is the identity relation.

- c knows that it is not safe for a to announce p in s , and thus it is not safe for c either: if she announces p in s , b would know that it was her since b knows that it is not safe for a to announce p in s .
- It is thus not safe for b to announce p in s either
- It is not safe for a or b to announce p in t (only two agents)
- p cannot be announced by anyone in u .

Thus, no safe announcements can be made in this model either.

These examples show that it is not enough that three agents know ϕ to ensure that ϕ can be safely announced by any of them; to be safe those three agents not only need to know ϕ but also that ϕ is safe, i.e., that three agents know ϕ and that ϕ is safe . . . and so on. This clearly has the flavour of *common knowledge*. However, the existence of a group of three agents having common knowledge of ϕ ,

$$M, s \models \bigvee_{a \neq b \neq c} C_{\{a,b,c\}} \phi$$

is sufficient but not necessarily for ϕ to be safe in s . A weaker condition would also be sufficient: we don't need the three agents to know that *the same* three agents safely can announce ϕ . It might be, for example, that a considers it possible that a , b and d safely can announce ϕ .

In order to define that weaker condition, let us recall the fixpoint definition of common knowledge (see, e.g., [17]). It is well known that $C_G \phi$ is the greatest

fixpoint of

$$E_G(\phi \wedge x)$$

w.r.t. a variable x not occurring in ϕ , or, more accurately, given an epistemic model M the extension $\|C_G\phi\|$ (the set of states where $C_G\phi$ is true) is the greatest fixpoint of the function

$$f(S) = \|E_G(\phi \wedge x)\|_{V[x=S]}^M$$

where $\|\psi\|_{V[x=S]}^M$ is the extension of ψ in the model M where the valuation function has been changed so that the extension of the variable x is S . Using a well known property of the fixed-point of monotonic functions, i.e., that the greatest fixed-point is equal to the union of all post-fixed points, we have that:

$$\|C_G\phi\| = \bigcup \left\{ S : S \subseteq \|E_G(\phi \wedge x)\|_{V[x=S]}^M \right\}.$$

In terms of the modal μ -calculus [12], this can be written explicitly:

$$C_G\phi \leftrightarrow \nu x. E_G(\phi \wedge x).$$

We now define a similar, weaker, notion of common knowledge in order to capture safety. Let us introduce a safety operator: $\blacktriangle\phi$ means that ϕ is safe in the current state. We let, where $N^3 = \{\{a, b, c\} : a, b, c \in N; a \neq b, a \neq c, b \neq c\}$ is the set of all groups of three different agents:

$$M, s \models \blacktriangle\phi \Leftrightarrow s \in \bigcup \left\{ S : S \subseteq \bigvee_{G \in N^3} E_G(\phi \wedge x) \right\}_{V[x=S]}^M.$$

In other words, $\blacktriangle\phi$ is the greatest fixpoint of $\bigvee_{G \in N^3} E_G(\phi \wedge x)$. The following shows that this definition makes sense, w.r.t. safety. It is similar to the iterative definition of common knowledge.

Lemma 1

$$\begin{aligned} M, s \models \blacktriangle\phi \text{ iff } & M, s \models \phi \text{ and} \\ & M, s \models \bigvee_{G_1 \in N^3} E_{G_1}\phi \text{ and} \\ & M, s \models \bigvee_{G_1 \in N^3} E_{G_1}(\phi \wedge \bigvee_{G_2 \in N^3} E_{G_2}\phi) \text{ and} \\ & M, s \models \bigvee_{G_1 \in N^3} E_{G_1}(\phi \wedge \bigvee_{G_2 \in N^3} E_{G_2}(\phi \wedge \bigvee_{G_3 \in N^3} E_{G_3}\phi)) \text{ and} \\ & \dots \end{aligned}$$

Lemma 1 also hints at another point of similarity between common knowledge and our new operator that the reader may find helpful. Recall that $C_G\phi$ is the reflexive transitive closure of \sim_a for all $a \in G$. Intuitively, this means that when interpreted on a model, $C_G\phi$ holds if and only if all paths labelled by any $a \in G$ only lead to ϕ -states. For safety, the difference is that now at each step of a path we choose a triple of agents such that *all* their relations lead to a ϕ -state. The triple of agents can be different at each step.

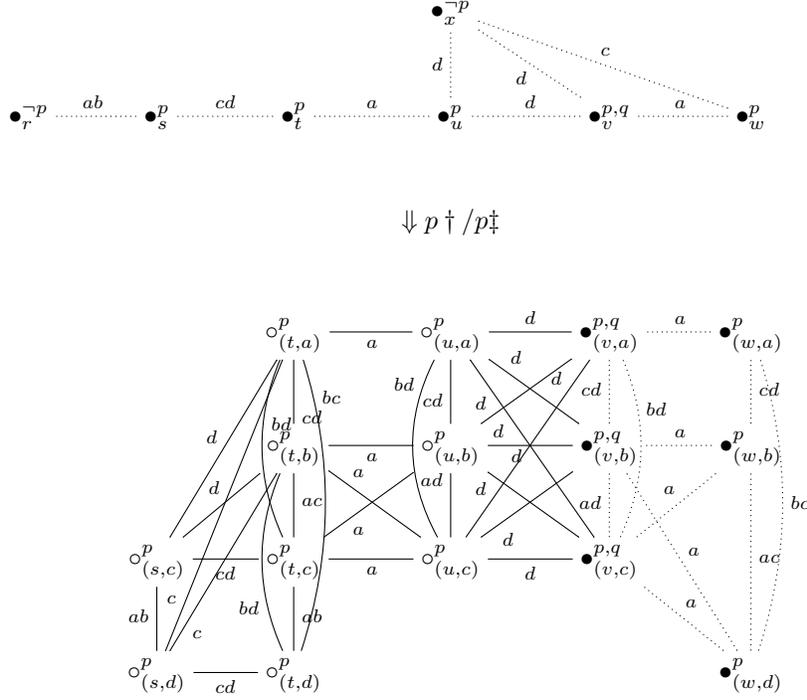


Figure 6: Four-agent epistemic model M , with updated model $M^{p\dagger}$ (black nodes and dotted edges) as well as $M^{p\ddagger}$ (black and white nodes, solid and dotted edges).

While $\blacktriangle\phi$ means that ϕ can safely be announced, $K_a\blacktriangle\phi$ means that ϕ can safely be announced by a . It is easy to see that:

$$\begin{aligned}
M, s \models K_a\blacktriangle\phi \text{ iff } & M, s \models \phi \text{ and} \\
& M, s \models \bigvee_{G_1 \in N^3, a \in G_1} E_{G_1}\phi \text{ and} \\
& M, s \models \bigvee_{G_1 \in N^3, a \in G_1} E_{G_1}(\phi \wedge \bigvee_{G_2 \in N^3} E_{G_2}\phi) \text{ and } \dots
\end{aligned}$$

Finally we can introduce an operator for intentional anonymous announcements, $[\phi\ddagger]$. The semantics is defined as follows. The model update is defined by strengthening the pre-condition in the definition of the updated model from the previous section from $K_a\phi$ to $K_a\blacktriangle\phi$ – a knows that ϕ is safe.

Definition 11 *The update of epistemic model $M = (S, \sim, V)$ by the safe pseudo-anonymous announcement of ϕ is the epistemic model $M^{\phi\ddagger} = (S', \sim', V')$ where:*

- $S' = \{(s, a) : s \in S, M, s \models K_a\blacktriangle\phi\}$
- $(s, a) \sim'_c (t, b)$ iff $s \sim_c t$ and $a = c$ iff $b = c$

- $V'(s, a) = V(s)$

We then let:

$$M, s \models [\phi_{\ddagger}]\psi \Leftrightarrow \forall a \in N, (M, s \models K_a \blacktriangle \phi \Rightarrow M^{\phi_{\ddagger}}, (s, a) \models \psi).$$

Thus, we have introduced two new operators: \blacktriangle and $[\phi_{\ddagger}]$. Do we want both in the formal language? Let's consider languages for all three combinations, which turns out to be useful later².

$$\begin{aligned} \mathcal{L}_{\ddagger} & \quad \phi ::= p \mid \neg\phi \mid \phi \wedge \phi \mid K_i \phi \mid [\phi_{\ddagger}]\phi \\ \mathcal{L}_{\blacktriangle} & \quad \phi ::= p \mid \neg\phi \mid \phi \wedge \phi \mid K_i \phi \mid \blacktriangle \phi \\ \mathcal{L}_{\ddagger\blacktriangle} & \quad \phi ::= p \mid \neg\phi \mid \phi \wedge \phi \mid K_i \phi \mid [\phi_{\ddagger}]\phi \mid \blacktriangle \phi \end{aligned}$$

Let's look at some examples. In the model in Figure 6:

- $M, s \models \neg\blacktriangle p$. p can't be safely announced in s : only two agents know p .
- $M, t \models \neg\blacktriangle p$. While all four agents know p in t , they don't all know that three agents know p : c and d consider it possible that only two agents know p . In other words, even though they know p they don't know that it can be safely announced, because they consider it possible that we are in s in which it can't be.
- $M, u \models \neg\blacktriangle p$. Three agents, a, b and c , know p . They also know that at least three agents know p . But they still don't know that p is safe: they don't know that at least three agents know that at least three agents know p .
- $M, v \models \blacktriangle p$ and $M, w \models \blacktriangle p$. p can be safely announced in v and in w . In fact, $\{w, v\}$ is the greatest fixed-point of $\bigvee_{G \in N^3} E_G(\phi \wedge x)$.
- $M, v \models K_a \blacktriangle p$. p can be safely announced by a in v .
- Even though a knows that p is safe in v , she doesn't know *why*. She considers it possible that we indeed are in v in which it would be safe for a, b and c to announce p , or in w in which it would be safe for a, b and d to announce p , but she doesn't know which of the two states we are in. But she knows that no matter which state we are in, it is safe to announce p – but for different reasons. We can say that she knows *de dicto* that p is safe, but not *de re*.
- $M, v \models [p_{\ddagger}]K_c K_d q$. After an intentionally anonymous announcement of p in w , c knows that d has learned that q is true.
- Safe anonymous announcements are *more informative* than general pseudo-anonymous announcements.

² \mathcal{L}_{\ddagger} is well defined without \blacktriangle in the syntax, even though \blacktriangle is used in the semantic condition for $[\phi_{\ddagger}]$. That semantic condition could of course be written without mentioning \blacktriangle , replacing it with its meaning.

- In every state in the updated model, *only the announcer knows who the announcer was.*

The latter point is no coincidence. The following shows that safe announcements are indeed safe:

Lemma 2 *In any update by a safe pseudo-anonymous announcement, it is common knowledge that no-one except the announcer knows who the announcer is (more technically: in any state (s, a) , for any agent $i \neq a$ there is a state (s', a') such that $(s, a) \sim_i (s', a')$ and $a \neq a'$).*

Proof: Let M be an epistemic model. Let (s, a) be a state in $M^{\phi\dagger}$ and let $d \neq a$. By definition, $M, s \models K_a \blacktriangle \phi$, so there are $b \neq c$, $b \neq a$ and $c \neq a$, such that $M, s \models K_b \blacktriangle \phi \wedge K_c \blacktriangle \phi$. That means that (s, b) and (s, c) are states in the updated model. We must have that either $b \neq d$ or $c \neq d$ (or both); assume the former. Since all of a , b and d are different, we have that $(s, a) \sim'_d (s, b)$. \square

Thus, $K_a \blacktriangle \phi$ is a *sufficient* condition for a safe anonymous announcement: after it is announced, no one will know who the announcer is. It is also *necessary*, in the sense that a must know that at least two other agents know ϕ (if there is only one "the other" agent will know who made the announcement), and must know that those two agents know that at least two other agents know, and so on $- K_a \blacktriangle \phi$. That justifies the definition of the model update: $K_a \blacktriangle \phi$ is the information that is revealed.

The reader might have observed that the updated model $M^{p\dagger}$ in Figure 6 is bisimilar to the submodel of M it is projected on, i.e., the submodel consisting of states v and w – exactly the submodel where $K_a \blacktriangle p$ holds for some a in other words. That is no coincidence: it is in fact always the case³. Indeed, we in fact have the following.

Lemma 3 *For any ϕ , ψ , M , and s ,*

$$M, s \models [\phi\dagger]\psi \text{ iff } M, s \models \blacktriangle \phi \Rightarrow M^{\blacktriangle \phi!}, s \models \psi$$

Proof: We first show that $M^{\phi\dagger}$ and $M^{\blacktriangle \phi!}$ are bisimilar. Let Z be such that $(s, a)Zs'$ iff $s = s'$ and let $(s, a)Zs$. *Atoms* is straightforward. *Forth* is as well: if $(s, a) \sim_c (t, b)$, then $s \sim_c t$ by definition. For *Back*, let $s \sim_c t$. First, consider the case that $c \neq a$. Since t is in $M^{\blacktriangle \phi!}$, $M, t \models \blacktriangle \phi$. Thus, there are $b, c, d \in N$ (all different) such that $M, t \models K_b \blacktriangle \phi \wedge K_c \blacktriangle \phi \wedge K_d \blacktriangle \phi$. At least one of b, c, d must be different from both a and c . Say, b . Thus $(s, a) \sim_c (t, b)$. Second, consider the case that $c = a$. Since $M, s \models K_c \blacktriangle$ and $s \sim_c t$, we also have that $M, t \models K_c \blacktriangle$. Thus (t, c) is in $M^{\phi\dagger}$, and we have that $(s, a) \sim_c (t, c)$.

Note that satisfaction of \blacktriangle is invariant under bisimulation (this follows, e.g., from Lemma 1). Then we have that $M, s \models [\phi\dagger]\psi$ iff for all a $M, s \models K_a \blacktriangle \phi$ implies $M^{\phi\dagger}, (s, a) \models \psi$ iff for all a $M, s \models K_a \blacktriangle \phi$ implies $M^{\blacktriangle \phi!}, s \models \psi$. We

³Unlike for $M^{\phi\dagger}$, see Figure 1.

argue that this is equivalent to $M, s \models \blacktriangle\phi$ implies $M^{\blacktriangle\phi!}, s \models \psi$. For the implication from left to right, if $M, s \models \blacktriangle\phi$ then $M, s \models K_a\blacktriangle\phi$ for some a by the fixed-point property of \blacktriangle (actually that holds for three different $a \in N$). From $M, s \models K_a\blacktriangle\phi \Rightarrow M^{\blacktriangle\phi!}, s \models \psi$ for all a , we get that $M^{\blacktriangle\phi!}, s \models \psi$. For the other direction, let $a \in N$. If $M, s \models K_a\blacktriangle\phi$, then $M, s \models \blacktriangle\phi$ by reflexivity, so $M^{\blacktriangle\phi!}, s \models \psi$. \square

As an immediate corollary of Lemma 3 we have that if we allow public announcement operators in the language, then the following holds:

$$\models [\phi_{\dagger}]\psi \leftrightarrow [\blacktriangle\phi!]\psi.$$

So, safe anonymous announcements are exactly public announcements of safety. Thus, we get a “simpler”, alternative, logically equivalent semantics. The existence of the original semantics and the fact that it corresponds exactly to this alternative semantics is of course still crucial: it is what ensures that safety is safe (Lemma 2)⁴.

It also means that PAL extended with \blacktriangle can express safe anonymous announcements. However, as we shall see when we now move on to comparing the expressive power of the three languages introduced above, even less is needed.

4.1 Expressive power

Consider the three languages $\mathcal{L}_{\dagger\blacktriangle}$, $\mathcal{L}_{\blacktriangle}$ and \mathcal{L}_{\dagger} . The first is similar to AMLC – it is epistemic logic extended with (restricted) action model operators and the common-knowledge-like operator \blacktriangle . In the same sense, $\mathcal{L}_{\blacktriangle}$ is similar to ELC. \mathcal{L}_{\dagger} is in-between, it would correspond to AML with common knowledge allowed in pre-conditions. AMLC is strictly more expressive than ELC [14, Chapter 8]; in particular expressions of the form $[M, s]C_G\phi$ cannot always be reduced to an ELC formula. Nevertheless, we shall now show that, perhaps contrary to intuition, *all three languages $\mathcal{L}_{\dagger\blacktriangle}$, $\mathcal{L}_{\blacktriangle}$ and \mathcal{L}_{\dagger} are actually all equally expressive*.

Let us start with the “intermediate” language \mathcal{L}_{\dagger} . It is easy to see that ϕ_{\dagger} can be expressed by action models like in Section 3.1, except that the preconditions now are $K_a\blacktriangle\phi$ instead of $K_a\phi$. We saw that \mathcal{L}_{\dagger} is reducible to epistemic logic, in the sense that for every formula in \mathcal{L}_{\dagger} there is an equivalent formula in epistemic logic, because we can translate the \mathcal{L}_{\dagger} formula to an AML formula which we know we can translate to an epistemic logic formula by using reduction axioms for AML. The same argument cannot be used directly for $\mathcal{L}_{\dagger\blacktriangle}$, because the preconditions, that the reduction axioms use, include the \blacktriangle operator, which

⁴We used invariance under standard bisimulation as a technical tool to establish the correspondence. Of course, when we say that “only the announcer knows who the announcer is” and so on, we implicitly mean agents can potentially discern between states (s, a) and (s, b) , even though, as shown above, they are bisimilar. An extended notion of bisimulation could be introduced to take into account this assumption, but it is in any case not picked up by the logical language, and here we only need standard bisimulation as a technical tool. It is of course also possible to think about ways this distinction *could* be picked up by the language, e.g., by introducing special atoms, but that is complicated by, e.g. the possibility of iterated announcements leading to states like $((s, a), b)$ and so on.

is not a part of the language. And if it was, as in the $\mathcal{L}_{\dagger\blacktriangle}$ language, we would need a reduction axiom for \blacktriangle as well. The corresponding result does, however, still hold.

Lemma 4 *For any $\phi \in \mathcal{L}_{\dagger}$, there is a $\phi' \in \mathcal{L}_{\blacktriangle}$ such that for any M, s , it holds that $M, s \models \phi$ iff $M, s \models \phi'$.*

Proof: Consider the language $AML\blacktriangle^-$, which is the language of AML extended with \blacktriangle , but with the restriction that \blacktriangle is not allowed in the scope of a $[M, s]$ modality. It is easy to see that the reduction axioms for AML still are valid for this extended language. Take the formula ϕ , let ψ be the corresponding AML formula. Now we recursively reduce the number of occurrences of $[\cdot]$ modalities in ψ by starting with an *outermost* one, i.e., one that is not in the scope of any other. Use the corresponding reduction axiom, and repeat until all $[\cdot]$ modalities are gone. Every time we use a reduction axiom we might introduce a new subformula $pre(s) = K_a\blacktriangle\psi'$, but never in the scope of a $[\cdot]$ modality since we started with the outermost modality. Thus, every time we use a reduction axiom the result is in $AML\blacktriangle^-$, and when we are done the result is in pure epistemic logic extended with \blacktriangle – which is $\mathcal{L}_{\blacktriangle}$. \square

Thus, \mathcal{L}_{\dagger} can be “reduced” to $\mathcal{L}_{\blacktriangle}$ (even though the latter is not a sub-language).

The second “surprise” is that the \blacktriangle operator can actually be expressed by $[\phi_{\dagger}^{\dagger}]$.

Lemma 5 *For any M, s and any ϕ (in any of the languages), $M, s \models \blacktriangle\phi$ iff $M, s \models \neg[\phi_{\dagger}^{\dagger}]\perp$.*

Proof: $M, s \models \neg[\phi_{\dagger}^{\dagger}]\perp$ iff there is an $a \in N$ such that $M, s \models K_a\blacktriangle\phi$ and $M^{\phi_{\dagger}^{\dagger}}, (s, a) \not\models \perp$ iff there is an $a \in N$ such that $M, s \models K_a\blacktriangle\phi$ iff (by reflexivity in one direction and the fixed-point definition of \blacktriangle in the other) $M, s \models \blacktriangle\phi$. \square

We thus get the mentioned result.

Corollary 2 $\mathcal{L}_{\dagger\blacktriangle} \approx \mathcal{L}_{\dagger} \approx \mathcal{L}_{\blacktriangle}$.

Proof: From Lemma 5 we get that $\mathcal{L}_{\dagger} \approx \mathcal{L}_{\dagger\blacktriangle}$. From Lemma 4 we get that $\mathcal{L}_{\dagger} \approx \mathcal{L}_{\blacktriangle}$, and from Lemma 5 again that $\mathcal{L}_{\blacktriangle} \approx \mathcal{L}_{\dagger}$. \square

Going back to the comparison with AMLC above and the mentioned non-reduction of $[M, s]C_G\psi$, the difference is that in $[\phi_{\dagger}^{\dagger}]\blacktriangle\psi$ we can replace $\blacktriangle\psi$ with $\neg[\psi_{\dagger}^{\dagger}]\perp$ and then reduce $[\phi_{\dagger}^{\dagger}]\neg[\psi_{\dagger}^{\dagger}]\perp$. This reduction might re-introduce \blacktriangle -operators, but outside the scope of any announcement modalities.

Thus, again the dynamic anonymous announcement operators can actually be expressed in a purely static language: $\mathcal{L}_{\blacktriangle}$. We thus move on to axiomatising safety.

all instances of propositional tautologies	Prop
From $\phi \rightarrow \psi$ and ϕ , derive ψ	Modus ponens
$K_a(\phi \rightarrow \psi) \rightarrow (K_a\phi \rightarrow K_a\psi)$	Distribution
$K_a\phi \rightarrow \phi$	Truth
$\neg K_a\phi \rightarrow K_a\neg K_a\phi$	Negative introspection
From ϕ , derive $K_a\phi$	Necessitation
$\blacktriangle\phi \rightarrow \bigvee_{G \in N^3} E_G(\phi \wedge \blacktriangle\phi)$	Mix
From $\phi \rightarrow \bigvee_{G \in N^3} E_G\phi$, derive $\phi \rightarrow \blacktriangle\phi$	Induction
From $\phi \rightarrow \psi$, derive $\blacktriangle\phi \rightarrow \blacktriangle\psi$	Monotonicity

Table 1: The proof system $\mathbf{S}\blacktriangle$.

4.2 Axiomatisation of safety

Observe that \blacktriangle does not distribute over implication: $\not\models \blacktriangle(\phi \rightarrow \psi) \rightarrow (\blacktriangle\phi \rightarrow \blacktriangle\psi)$. In particular, $\bigvee_{G \in N^3}(\phi \rightarrow \psi)$ and $\bigvee_{G \in N^3}\phi$ does not imply $\bigvee_{G \in N^3}\psi$ – it might not be the same G in the first two cases. For similar reasons the conjunctive closure axiom, $(\blacktriangle\phi \wedge \blacktriangle\psi) \rightarrow \blacktriangle(\phi \wedge \psi)$ does not hold. This is similar to *somebody knows* [5], but unlike most other group knowledge operators which are normal modalities. The other direction of conjunctive closure, *monotonicity*, $\blacktriangle(\phi \wedge \psi) \rightarrow (\blacktriangle\phi \wedge \blacktriangle\psi)$, does hold (again, similarly to somebody knows).

The axiomatic system $\mathbf{S}\blacktriangle$ for $\mathcal{L}_\blacktriangle$ is shown in Table 1. The first two parts is a standard axiomatisation of propositional logic and the individual knowledge operators. The Monotonicity rule combines the monotonicity axiom and the replacement of equivalents rule standard in weak modal logics. We note that necessitation for \blacktriangle , from ϕ derive $\blacktriangle\phi$, follows (we have $\blacktriangle\top$ from the Induction rule). The Mix axiom (or *fixed-point axiom*) says that $\blacktriangle\phi$ is indeed a fixed-point of $\bigvee_{G \in N^3} E_G(\phi \wedge x)$. Finally, the Induction rule give us a way to derive $\blacktriangle\phi$. Mix and the Induction rule can be seen as adaptations of similar axioms/rules for common knowledge, see, e.g., [17, 14].

We will use the following shorthand:

$$\blacktriangle_n = \phi \wedge \bigvee_{G_1 \in N^3} E_{G_1}(\phi \wedge \bigvee_{G_2 \in N^3} E_{G_2}(\phi \wedge \bigvee_{G_3 \in N^3} E_{G_3}(\phi \wedge \dots \wedge \bigvee_{G_n \in N^3} E_{G_n}\phi))).$$

Thus, $\blacktriangle_0 = \phi$, $\blacktriangle_1 = \phi \wedge \bigvee_{G_1 \in N^3} E_{G_1}\phi$, and so on, and

$$M, s \models \blacktriangle\phi \text{ iff } M, s \models \blacktriangle_n\phi \text{ for all } n. \quad (1)$$

4.2.1 Soundness

Lemma 6 (Mix Axiom) $\models \blacktriangle\phi \rightarrow \bigvee_{G \in N^3} E_G(\phi \wedge \blacktriangle\phi)$

Proof: By definition, $\|\blacktriangle\phi\|$ is a fixed-point of the function $f_{\bigvee_{G \in N^3}(\phi \wedge x)}(A)$ that maps a set of states A to the set of states where $\bigvee_{G \in N^3}(\phi \wedge x)$ is true

when x is true exactly in A . It follows immediately that $M, s \models \blacktriangle\phi$ iff $M, s \models \bigvee_{G \in N^3} E_G(\phi \wedge \blacktriangle\phi)$. \square

Lemma 7 (Induction Rule) *If $\models \phi \rightarrow \bigvee_{G \in N^3} E_G\phi$, then $\models \phi \rightarrow \blacktriangle\phi$.*

Proof: Let $\models \phi \rightarrow \bigvee_{G \in N^3} E_G\phi$ and let $M, s \models \phi$. We show that $M, s \models \blacktriangle_n$ for all n , by induction on n . The base case, $n = 0$, is immediate. Assume that $M, s \models \blacktriangle_n$, i.e., that

$$M, s \models \phi \wedge \bigvee_{G_1 \in N^3} E_{G_1}(\phi \wedge \bigvee_{G_2 \in N^3} E_{G_2}(\phi \wedge \bigvee_{G_3 \in N^3} E_{G_3}(\phi \wedge \dots \wedge \bigvee_{G_n \in N^3} E_{G_n}\phi))).$$

It follows from validity of $\phi \rightarrow \bigvee_{G \in N^3} E_G\phi$ that also

$$M, s \models \phi \wedge \bigvee_{G_1 \in N^3} E_{G_1}(\phi \wedge \dots \wedge \bigvee_{G_n \in N^3} E_{G_n}(\phi \wedge \bigvee_{G_{n+1} \in N^3} E_{G_{n+1}}\phi)),$$

i.e., $M, s \models \blacktriangle_{n+1}\phi$. \square

We thus get the following.

Proposition 3 $\mathbf{S}\blacktriangle$ is sound.

Proof: Validity (preservation) for the two first parts of the system is standard. Monotonicity is straightforward. Validity of the Mix axiom and validity preservation for the induction rule were shown in Lemmas 6 and 7, respectively. \square

4.2.2 Completeness

We now prove that $\mathbf{S}\blacktriangle$ is also complete. Similar to epistemic logic with common knowledge, it is easy to see that the logic is not *compact*. For example,

$$\{\blacktriangle_n\phi : n \geq 0\} \cup \{\neg\blacktriangle\phi\}$$

is not satisfiable, but any finite subset of it is. Thus, we have to settle for *weak* rather than *strong* completeness. To that end, we will adapt the standard technique for common knowledge of defining a finite canonical model based on the finite syntactic closure of some formula. There is a complication: unlike C_G , \blacktriangle is not a normal modality. This shows up in the axiomatisation: the K axiom is replaced with the (weaker) monotonicity rule. The consequence, aside from the fact that we cannot rely on distribution over implication like in proofs for common knowledge, is that there is no normal relational semantics for \blacktriangle (like “ ϕ is true on all states reachable by a G -path” for $C_G\phi$). We now define an alternative definition of the semantics and show that it is equivalent, that we will use in the completeness proof.

Given a model M , a *group assignment function* for M is a function $f : W \rightarrow \mathcal{P}(N)$ such that $f(s) \in N^3$ or $f(s) = \emptyset$ for all s , i.e., assigning a group of three

agents to some of the states, such that for all $s, t \in W$ and $i \in f(s)$, if $s \sim_i t$ then $f(t) \neq \emptyset$. The idea is that f works as an assignment of groups of three agents to certain states, such that if we follow the accessibility relations for those agents, f again assigns a group of three agents to those states, and so on.

An f -consistent path in M is a finite sequence of states $s_0 s_1 s_2 \cdots s_m$ such that for all $0 \leq i < m$, $s_i \sim_a s_{i+1}$ for some $a \in f(s_i)$. We say that ϕ is true on a path if it is true in every state on that path.

Lemma 8 $M, s \models \blacktriangle \phi$ iff there is a group assignment function f such that $f(s) \neq \emptyset$ and ϕ is true on every f -consistent path starting in s .

Proof: *Left-to-right:* let $M, s \models \blacktriangle \phi$. Define f as follows: for any state t , $f(t) = G$ for some G such that $M, t \models E_G \blacktriangle \phi$ if such a G exists, $f(t) = \emptyset$ if not. If there are several such G chose any of them. We must show that f is indeed a group assignment function. Let $i \in f(s')$ and $s' \sim_i t$; we must show that $f(t) \neq \emptyset$. Since $i \in f(s')$, there is a $G \subseteq N$ such that $M, s' \models E_G \blacktriangle \phi$ and $i \in G$. Thus $M, t \models \blacktriangle \phi$. By the Mix axiom, $M, t \models E_{G'} \blacktriangle \phi$ for some G' , so $f(t) \neq \emptyset$. Finally, $f(s) \neq \emptyset$ from $M, s \models \blacktriangle \phi$ and validity of Mix.

We proceed by induction on the length of any f -consistent path starting in s . The base case is immediate. Consider an f -consistent path $s_0 \cdots s_n s_{n+1}$ where $s_0 = s$. By the induction hypothesis s_0, \dots, s_n are ϕ -states. Since $s_n \sim_a s_{n+1}$ for some $a \in f(s_n)$ and $M, s_n \models E_{f(s_n)} \blacktriangle \phi$, also $M, s_{n+1} \models \phi$.

Right-to-left: we reason by contraposition. Assume that $M, s \not\models \blacktriangle \phi$, i.e., that $M, s \not\models \blacktriangle_n \phi$ for some n . Let f be a group assignment function such that $f(s) \neq \emptyset$. We must show that there is a f -consistent path starting in s where ϕ is not true. We have that:

$$M, s \models \neg \phi \vee \bigwedge_{G_1 \in N^3} \bigvee_{i_1 \in G_1} \hat{K}_{i_1} (\neg \phi \vee \bigwedge_{G_2 \in N^3} \bigvee_{i_2 \in G_2} \hat{K}_{i_2} (\neg \phi \vee \cdots \vee \bigwedge_{G_n \in N^3} \bigvee_{i_n \in G_n} \hat{K}_{i_n} \neg \phi)).$$

Thus, there is a $k \leq n$ such that

$$M, s \models \bigwedge_{G_1 \in N^3} \bigvee_{i_1 \in G_1} \hat{K}_{i_1} \bigwedge_{G_2 \in N^3} \bigvee_{i_2 \in G_2} \hat{K}_{i_2} \bigwedge_{G_3 \in N^3} \bigvee_{i_3 \in G_3} \hat{K}_{i_3} \cdots \vee \bigwedge_{G_k \in N^3} \bigvee_{i_k \in G_k} \hat{K}_{i_k} \neg \phi.$$

Thus there is a path $s_0 s_1 \cdots s_{k+1}$ where:

$$\begin{array}{ll} s_0 = s & \\ s_0 \sim_{i_1} s_1 & \text{some } i_1 \in f(s_0) \\ s_1 \sim_{i_2} s_2 & \text{some } i_2 \in f(s_1) \\ \dots & \\ s_k \sim_{i_{k+1}} s_{k+1} & \text{some } i_{k+1} \in f(s_k) \end{array}$$

such that $M, s_{k+1} \models \neg \phi$. This is an f -consistent path. \square

We now proceed with defining the finite canonical model, and proving a truth lemma. We adapt the standard proof for common knowledge (see [14,

Chapter 7]⁵).

Definition 12 *The closure $cl(\gamma)$ of a formula γ is the smallest set that contains all subformulas of γ , is closed under single negations, and that contains $E_G \blacktriangle \phi$ for all $G \subseteq N$ whenever it contains $\blacktriangle \phi$.*

The closure of any formula is finite. As usual we say that a set of formulas Γ is maximal consistent in $cl(\gamma)$ for some given γ iff $\Gamma \subseteq cl(\gamma)$, Γ is consistent, and there is no consistent $\Gamma' \subseteq cl(\gamma)$ such that $\Gamma \subset \Gamma'$. When Δ is a set of formulas we write $\underline{\Delta}$ for $\bigwedge_{\delta \in \Delta} \delta$.

The canonical model is defined as follows.

Definition 13 *The canonical model for some formula γ is $M^\gamma = (S^\gamma, \sim^\gamma, V^\gamma)$ where:*

- S^γ is the set of all sets of formulas maximal consistent in $cl(\gamma)$,
- $\Gamma \sim^\gamma \Delta$ iff $\{K_i \phi : K_i \phi \in \Gamma\} = \{K_i \phi : K_i \phi \in \Delta\}$, and
- $V^\gamma(p) = \{\Gamma \in S^\gamma : p \in \Gamma\}$.

This definition of the canonical model is identical to the one used for common knowledge [14], except that the closure is wider (it contains $K_i \blacktriangle \phi$ for every i , instead of $K_i C_G \phi$ for every $i \in G$). Many of the properties of that model, like deductive closure of maximal consistent sets in $cl(\gamma)$, carry over, and we only need to focus on the following property for the \blacktriangle modality.

A ϕ -path in the canonical model is a path $\Gamma_0 \Gamma_1 \dots$ where $\phi \in \Gamma_i$ for every i .

Lemma 9 *For any formula γ and $\Gamma \in S^\gamma$, if $\blacktriangle \phi \in cl(\gamma)$ then $\blacktriangle \phi \in \Gamma$ iff there is a group assignment function f such that $f(\Gamma) \neq \emptyset$ and every f -consistent path from Γ is a ϕ -path.*

Proof: *Left-to-right:* let $\blacktriangle \phi \in \Gamma$. We define f as follows: for any $\Delta \in S^\gamma$, $f(\Delta) = G$ if $E_G \blacktriangle \phi \in \Delta$ for some $G \subseteq N$ and $f(\Delta) = \emptyset$ otherwise. If there are several such G , chose one of them. Consider a Δ such that $f(\Delta) \neq \emptyset$ and $\Delta \sim_i^\gamma \Delta'$ for some $i \in f(\Delta)$. We must show that $f(\Delta') \neq \emptyset$. Since $K_i \blacktriangle \phi \in \Delta$, $K_i \blacktriangle \phi \in \Delta'$ by definition of \sim_i^γ , and thus $\blacktriangle \phi \in \Delta'$ by the Truth axiom, and $E_{G'} \blacktriangle \phi \in \Delta'$ for some G' by Mix and the fact that $E_{G'} \blacktriangle \phi \in cl(\gamma)$. Finally, $f(\Gamma) \neq \emptyset$ also follows from Mix and the definition of the closure.

We must show that every f -consistent path from Γ is a ϕ -path. The proof is by induction on the length of the path. For the base case, we have that $\phi \in \Gamma$ from the Mix and Truth axioms. For the induction step, consider an f -consistent path $\Gamma_0 \dots \Gamma_n \Gamma_{n+1}$ where $\phi \in \Gamma_j$ for all $j \leq n$. From f -consistency

⁵In addition to the complications mentioned above, there are other differences, including that the correspondent to the *induction axiom* for common knowledge, $\blacktriangle(\phi \rightarrow \bigvee_{G \in N^3} \phi) \rightarrow (\phi \rightarrow \blacktriangle \phi)$, does not hold. We use a variant of the induction rule for common knowledge, used in, e.g., [17]

we have that $\Gamma_n \sim_i^\gamma \Gamma_{n+1}$ for some $i \in G$ for some G such that $E_G \blacktriangle \phi \in \Gamma_n$. It follows from the definition of \sim_i^γ that $K_i \blacktriangle \phi \in \Gamma_{n+1}$, and thus that $\blacktriangle \phi \in \Gamma_{n+1}$ from the Truth axiom, and thus that $\phi \in \Gamma_{n+1}$ from the Mix axiom and Truth axiom again.

Right-to-left: let f be such that $f(\Gamma) \neq \emptyset$ and every f -consistent path is a ϕ -path. Let $S_{f,\phi}$ be the set of all sets Δ maximal consistent in $cl(\gamma)$ such that $f(\Delta) \neq \emptyset$ and every f -consistent path from Δ is a ϕ -path. Let $\chi = \bigvee_{\Delta \in S_{f,\phi}} \underline{\Delta}$. We first show that

$$\vdash \chi \rightarrow \bigvee_{G \in N^3} E_G \chi \quad (2)$$

Assume, towards a contradiction, that $\chi \wedge \neg \bigvee_{G \in N^3} E_G \chi$ is consistent. Then $\underline{\Delta} \wedge \neg \bigvee_{G \in N^3} E_G \chi$ is consistent for some $\Delta \in S_{f,\phi}$. Then, for any $G \in N^3$ there must be a $i_G \in G$ such that $\underline{\Delta} \wedge \hat{K}_{i_G} \neg \chi$ is consistent. $\underline{\Delta} \wedge \hat{K}_{i_G} \bigvee_{\Theta \in S^\gamma \setminus S_{f,\phi}} \underline{\Theta}$ is consistent for each i_G , and by modal reasoning for individual knowledge, $\underline{\Delta} \wedge \bigvee_{\Theta \in S^\gamma \setminus S_{f,\phi}} \hat{K}_{i_G} \underline{\Theta}$ is consistent for each i_G . Thus, for every $G \in N^3$ there is an $i_G \in G$ and a $\Theta_{i_G} \in S^\gamma \setminus S_{f,\phi}$ such that $\underline{\Delta} \wedge \hat{K}_{i_G} \underline{\Theta}_{i_G}$ is consistent. In particular, since $\Delta \in S_{f,\phi}$ and thus $f(\Delta) \neq \emptyset$, there is an $i_G \in G = f(\Delta)$ and a $\Theta_{i_G} \in S^\gamma \setminus S_{f,\phi}$ such that $\underline{\Delta} \wedge \hat{K}_{i_G} \underline{\Theta}_{i_G}$ is consistent. It follows (by a standard property of the canonical model [14, Item 4 of Lemma 7.14]) that $\Delta \sim_{i_G}^\gamma \Theta_{i_G}$. Since $i_G \in f(\Delta)$, $f(\Theta_{i_G}) \neq \emptyset$ and since Θ_{i_G} is not in $S_{f,\phi}$ that means that there is an f -consistent path from Θ_{i_G} which is not a ϕ -path. It follows that there is an f -consistent path from Δ that is not a ϕ -path. But $\Delta \in S_{f,\phi}$, which leads to a contradiction. Thus, we have shown (2).

From (2) and the induction rule, we get that $\vdash \chi \rightarrow \blacktriangle \chi$. Since Γ is one of the disjuncts in χ we have that $\vdash \underline{\Gamma} \rightarrow \chi$. Thus, $\vdash \underline{\Gamma} \rightarrow \blacktriangle \chi$, and $\blacktriangle \chi \in \Gamma$. Since $\phi \in \bigcap_{\Delta \in S_{f,\phi}} \Delta$, we also have that $\vdash \chi \rightarrow \phi$. By Monotonicity, $\vdash \blacktriangle \chi \rightarrow \blacktriangle \phi$, and since $\blacktriangle \phi \in cl(\gamma)$, we also have that $\blacktriangle \phi \in \Gamma$. \square

Lemma 10 (Truth) *For any formulas γ and ϕ and any $\Gamma \in S^\gamma$, $\phi \in \Gamma$ iff $M^\gamma, \Gamma \models \phi$.*

Proof: The proof is by induction on the structure of ϕ . We only show the induction step for $\phi = \blacktriangle \psi$, the other cases are straightforward and/or standard.

We have that $M^\gamma, \Gamma \models \blacktriangle \psi$ iff (by Lemma 8) there is an f such that $f(\Gamma) \neq \emptyset$ and ψ is true on every f -consistent path starting in Γ iff (by the induction hypothesis) there is an f such that $f(\Gamma) \neq \emptyset$ and every f -consistent path starting in Γ is a ψ -path iff (by Lemma 9) $\blacktriangle \psi \in \Gamma$. \square

We immediately get the following.

Theorem 1 (Completeness) S_\blacktriangle is complete.

Proof: If $\neg \phi$ is consistent it can be extended to a set Γ that is maximal consistent in $cl(\neg \phi)$. From the truth lemma we get that $M^{-\phi}, \Gamma \models \neg \phi$. \square

5 Discussion

Pseudo-anonymous public announcements $\phi\dagger$ are made by an agent inside the system, but not explicitly "signed" by that agent. In a sense, such announcements are in-between public announcements $\phi!$ and $K_a\phi!$. These announcements are pseudo-anonymous, because, depending on the background knowledge of the agents in a system, the identity of the announcer might still be revealed. We showed that like normal public announcements, these operators can actually be expressed in (pure) epistemic logic.

Similarly to the *Russian Cards Problem* [13] (see below), if we make assumptions about the *intentions* of the announcer, we *learn more*. Intentional anonymous announcements $\phi\dagger$ capture exactly the announcements that are guaranteed to ensure anonymity. It all boils down to *safety*: intentional anonymous announcements $\phi\dagger$ are exactly public announcements of safety $\blacktriangle\phi!$. Furthermore, the $\phi\dagger$ operators can be expressed in epistemic logic extended with only the safety operator \blacktriangle . We gave a complete axiomatisation of the latter.

There are still many open problems. While the expressivity results "reduce" the languages with $\phi\dagger$ and $\phi\dagger$ to equivalent languages we have axiomatisations for, they don't directly give us "native" axiomatisations of those languages themselves. For the case of pseudo-anonymous announcements, we don't get reduction axioms directly from AML, since the language is less expressive. We note that we would get reduction axioms (and thus completeness) immediately if we used a more fine-grained variant

$$[\phi\dagger_a]\psi$$

meaning "after ϕ is pseudo-anonymously announced by a (" a writes ϕ on the blackboard"), ψ is true" – corresponding to a single action model point (instead of a union) – with the following semantics:

$$M, s \models [\phi\dagger_a]\psi \Leftrightarrow [M, s \models K_a\phi \Rightarrow M^{\phi\dagger}, (s, a) \models \psi].$$

This would be completely deterministic, corresponding to exactly one pointed action model instead of the union, and like for general action models and union we would have the following:

$$M, s \models [\phi\dagger]\psi \text{ iff } M, s \models \bigwedge_{a \in N} [\phi\dagger_a]\psi.$$

A conceptually closely related work is van Ditmarsch's analysis of the *Russian Cards Problem* [13] which also models *safe announcements* – albeit with another notion of safety, namely that a secret "card deal" is not revealed instead of not revealing the identity of the announcer. A safe announcement of ϕ by a is captured by $[K_a\phi \wedge [K_a\phi]C\text{ignorant}]!$, where C is common knowledge among all agents and *ignorant* is the safety or secrecy condition – corresponding to only the announcer knowing the identity of the announcer in our case. That is again equivalent to the sequence $[K_a\phi!][C\text{ignorant}]!$, which means that ϕ is safe to

announce for a if $C\text{ignorant}$ holds after the announcement of $K_a\phi$. Our $[\phi\dagger]$ operator also satisfies that definition of safety (Lemma 2). This raises the question whether a safe anonymous announcement is equivalent to a pseudo-anonymous announcement followed by an announcement that it is common knowledge that no one except the announcer now knows her identity. We can't refer directly to the identity of "the announcer" in the updated model in the syntax, but even if we could, the answer to the question would be "no". Let onlyone mean that only the announcer now knows her identity. To see that

$$[\phi\dagger]\psi \leftrightarrow [\phi\dagger][C\text{onlyone!}]\psi$$

does not hold, consider Figure 6: the $\phi\dagger$ update has several more states and $C\text{onlyone}$ would not hold in any of them. One could perhaps consider iterations of the form $[\phi\dagger][\text{onlyone!}][\text{onlyone!}]\dots$, but that is also not equivalent to $[\phi\dagger]$ (for any number of iterations): as a counterexample see Figure 5. The first announcement of onlyone would remove the state (t, a) , since agent b knows in that state the identity of the announcer. The second announcement would remove the state (t, b) , but then it stops, resulting in a three-state model. The update with safe announcements however, fails as argued earlier – $\blacktriangle\phi$ does not hold in any state in the original model in Figure 5.

Regarding safety, we can potentially go even further by using some variant of *distributed knowledge* [21]. That would allow us to reason about anonymous announcements that are safe even when non-announcing agents share their knowledge to deduce the identity of the announcer.

Another avenue of further research is to extend the presented formalisms with *quantification over anonymous announcements*. This will allow us to express properties like "there is a safe anonymous announcement, such that ϕ holds afterwards". Quantification over standard public announcements is relatively well-studied [6, 3, 8, 7, 1, 2, 4], both in the presence of common and distributed knowledge, and we expect many of the existing intuitions hold also in the case of the anonymous announcements.

Also of interest for future work, is to model *self-referential* intentionally anonymous announcements of the form "after this very announcement, no-one will know", as studied recently in [9, 24]⁶.

Acknowledgments This work was supported by ROIS-DS-JOINT (047RP2024).

References

- [1] Ågotnes, T., Ditmarsch, H., French, T.: The undecidability of quantified announcements. *Studia Logica* **104**(4), 597–640 (2016). <https://doi.org/10.1007/s11225-016-9657-0>, <http://dx.doi.org/10.1007/s11225-016-9657-0>

⁶That also applies to the Russian Cards problem.

- [2] Ågotnes, T., Alechina, N., Galimullin, R.: Logics with group announcements and distributed knowledge: Completeness and expressive power. *Journal of Logic, Language and Information* **31**(2), 141–166 (2022). <https://doi.org/10.1007/s10849-022-09355-0>
- [3] Ågotnes, T., Balbiani, P., van Ditmarsch, H., Seban, P.: Group announcement logic. *Journal of Applied Logic* **8**(1), 62–81 (2010)
- [4] Ågotnes, T., Galimullin, R.: Quantifying over information change with common knowledge. *Autonomous Agents and Multi-agent Systems* **37**(19) (2023)
- [5] Ågotnes, T., Wáng, Y.N.: Somebody knows. In: Bienvenu, M., Lakemeyer, G., Erdem, E. (eds.) *Proceedings of the 18th KR*. pp. 2–11 (2021). <https://doi.org/10.24963/kr.2021/1>, <https://doi.org/10.24963/kr.2021/1>
- [6] Balbiani, P., Baltag, A., van Ditmarsch, H., Herzig, A., Hoshi, T., Lima, T.D.: ‘Knowable’ as ‘known after an announcement’. *Review of Symbolic Logic* **1**(3), 305–334 (2008)
- [7] Balbiani, P.: Putting right the wording and the proof of the truth lemma for apal. *Journal of Applied Non-Classical Logics* **25**(1), 2–19 (2015)
- [8] Balbiani, P., van Ditmarsch, H.: A simple proof of the completeness of apal. *Studies in Logic* **8**, 65–78 (2015)
- [9] Baltag, A., Bezhaniashvili, N., Fernández-Duque, D.: The topology of surprise. In: Kern-Isberner, G., Lakemeyer, G., Meyer, T. (eds.) *Proceedings of the 19th KR* (2022). <https://doi.org/10.24963/kr.2022/4>
- [10] Baltag, A., Moss, L.S.: Logics for epistemic programs. *Synthese* **139**(2), 165–224 (2004). <https://doi.org/10.1023/B:SYNT.0000024912.56773.5E>
- [11] Barthe, G., Gaboardi, M., Arias, E.J.G., Hsu, J., Kunz, C., Strub, P.: Proving differential privacy in hoare logic. In: *Proceedings of the 27th CSF*. pp. 411–424. IEEE Computer Society (2014). <https://doi.org/10.1109/CSF.2014.36>
- [12] Bradfield, J., Stirling, C.: Modal mu-calculi. In: Blackburn, P., Van Benthem, J., Wolter, F. (eds.) *Handbook of Modal Logic, Studies in Logic and Practical Reasoning*, vol. 3, pp. 721–756. Elsevier (2007). [https://doi.org/10.1016/S1570-2464\(07\)80015-2](https://doi.org/10.1016/S1570-2464(07)80015-2)
- [13] van Ditmarsch, H.: The russian cards problem. *Studia logica* **75**, 31–62 (2003). <https://doi.org/10.1023/A:1026168632319>
- [14] van Ditmarsch, H., van der Hoek, W., Kooi, B.: *Dynamic Epistemic Logic*, Synthese Library, vol. 337. Springer (2007)

- [15] van Ditmarsch, H., van der Hoek, W., Kooi, B., Kuijer, L.B.: Arrow update synthesis. *Information and Computation* **275**, 104544 (2020). <https://doi.org/10.1016/J.IC.2020.104544>
- [16] Domingo-Ferrer, J., Sánchez, D., Soria-Comas, J.: Database Anonymization: Privacy Models, Data Utility, and Microaggregation-based Inter-model Connections. *Synthesis Lectures on Information Security, Privacy, & Trust*, Morgan & Claypool Publishers (2016). <https://doi.org/10.2200/S00690ED1V01Y201512SPT015>
- [17] Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.Y.: Reasoning About Knowledge. The MIT Press, Cambridge, Massachusetts (1995)
- [18] Goranko, V., Otto, M.: Model theory of modal logic. In: Blackburn, P., van Benthem, J., Wolter, F. (eds.) *Handbook of Modal Logic*, Studies in Logic and Practical Reasoning, vol. 3, pp. 249–329. Elsevier (2007)
- [19] Halpern, J.Y., O’Neill, K.R.: Anonymity and information hiding in multiagent systems. *Journal of Computer Security* **13**(3), 483–512 (2005). <https://doi.org/10.3233/JCS-2005-13305>
- [20] Halpern, J.Y., O’Neill, K.R.: Secrecy in multiagent systems. *ACM Transactions on Information and System Security* **12**(1), 5:1–5:47 (2008). <https://doi.org/10.1145/1410234.1410239>
- [21] van der Hoek, W., van Linder, B., Meyer, J.J.: Group knowledge is not always distributed (neither is it always implicit). *Mathematical Social Sciences* **38**(2), 215–240 (1999). [https://doi.org/https://doi.org/10.1016/S0165-4896\(99\)00013-X](https://doi.org/https://doi.org/10.1016/S0165-4896(99)00013-X)
- [22] Hughes, D.J.D., Shmatikov, V.: Information hiding, anonymity and privacy: a modular approach. *Journal of Computer Security* **12**(1), 3–36 (2004). <https://doi.org/10.3233/JCS-2004-12102>
- [23] Jiang, J., Naumov, P.: De re/de dicto distinction: a logicians’ perspective on data anonymity. *Journal of Cybersecurity* **11**(1) (2025). <https://doi.org/10.1093/CYBSEC/TYAE025>
- [24] Li, Y., Ren, J., Ågotnes, T.: The surprise exam in full modal fixed-point logic. In: *Proceedings of the The 6th International Conference on Logic and Argumentation (CLAR 2025)*. LNCS, Springer (2025), to appear.
- [25] Machanavajjhala, A., Kifer, D., Gehrke, J., Venkatasubramanian, M.: *L*-diversity: Privacy beyond *k*-anonymity. *ACM Transactions on Knowledge Discovery from Data* **1**(1), 3 (2007). <https://doi.org/10.1145/1217299.1217302>
- [26] Mano, K., Kawabe, Y., Sakurada, H., Tsukada, Y.: Role interchange for anonymity and privacy of voting. *Journal of Logic and Computation* **20**(6), 1251–1288 (2010). <https://doi.org/10.1093/LOGCOM/EXQ013>

- [27] Meyer, J.C., van der Hoek, W.: Epistemic logic for AI and computer science, Cambridge tracts in theoretical computer science, vol. 41. CUP (1995)
- [28] More, S.M., Naumov, P.: Hypergraphs of multiparty secrets. *Annals of Mathematics and Artificial Intelligence* **62**(1-2), 79–101 (2011). <https://doi.org/10.1007/S10472-011-9252-Z>
- [29] More, S.M., Naumov, P.: Logic of secrets in collaboration networks. *Annals of Pure and Applied Logic* **162**(12), 959–969 (2011). <https://doi.org/10.1016/J.APAL.2011.06.001>
- [30] Naumov, P., Orejola, O.: Shhh! the logic of clandestine operations. In: *Proceedings of the 32nd IJCAI*. pp. 3304–3311. ijcai.org (2023). <https://doi.org/10.24963/IJCAI.2023/368>
- [31] Pfitzmann, A., Hansen, M.: A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management (Aug 2010), http://dud.inf.tu-dresden.de/literatur/Anon_Terminology_v0.34.pdf, v0.34
- [32] Plaza, J.: Logics of public communications. In: *Proceedings of the 4th ISMIS*. pp. 201–216. Oak Ridge National Laboratory (1989)
- [33] Schneider, S.A., Sidiropoulos, A.: CSP and anonymity. In: Bertino, E., Kurth, H., Martella, G., Montolivo, E. (eds.) *Proceedings of the 4th ESORICS*. LNCS, vol. 1146, pp. 198–218. Springer (1996). https://doi.org/10.1007/3-540-61770-1_38
- [34] Sweeney, L.: k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* **10**(5), 557–570 (2002). <https://doi.org/10.1142/S0218488502001648>
- [35] Syverson, P.F., Stubblebine, S.G.: Group principals and the formalization of anonymity. In: Wing, J.M., Woodcock, J., Davies, J. (eds.) *Proceedings of the FM 1999*. LNCS, vol. 1708, pp. 814–833. Springer (1999). https://doi.org/10.1007/3-540-48119-2_45
- [36] Tsukada, Y., Mano, K., Sakurada, H., Kawabe, Y.: Anonymity, privacy, onymity, and identity: A modal logic approach. In: *Proceedings of the 12th CSE*. pp. 42–51. IEEE Computer Society (2009). <https://doi.org/10.1109/CSE.2009.251>
- [37] Tsukada, Y., Sakurada, H., Mano, K., Manabe, Y.: On compositional reasoning about anonymity and privacy in epistemic logic. *Annals of Mathematics and Artificial Intelligence* **78**(2), 101–129 (2016). <https://doi.org/10.1007/S10472-016-9516-8>, <https://doi.org/10.1007/s10472-016-9516-8>

- [38] van Eijck, J., Orzan, S.: Epistemic verification of anonymity. *Electronic Notes in Theoretical Computer Science* **168**, 159–174 (2007). <https://doi.org/https://doi.org/10.1016/j.entcs.2006.08.026>, proceedings of the 2nd VODCA
- [39] Xiong, Z., Ågotnes, T.: The logic of secrets and the interpolation rule. *Annals of Mathematics and Artificial Intelligence* **91**(4), 375–407 (2023). <https://doi.org/10.1007/S10472-022-09815-0>
- [40] Ye, X., Li, Z., Li, Y.: Capture inference attacks for k-anonymity with privacy inference logic. In: Ramamohanarao, K., Krishna, P.R., Mohania, M.K., Nantajeewarawat, E. (eds.) *Proceedings of the 12th DASFAA*. LNCS, vol. 4443, pp. 676–687. Springer (2007). https://doi.org/10.1007/978-3-540-71703-4_57