Proceedings of Machine Learning Research vol 99:1–35, 2019 32nd Annual Conference on Learning Theory

Reasoning in Bayesian Opinion Exchange Networks Is PSPACE-Hard

Jan Hązła JHAZLA@MIT.EDU
Department of Mathematics; Institute for Data, Systems and Society; Massachusetts Institute of Technology

Ali Jadbabaie JADBABAI@MIT.EDU
Department of Civil and Environmental Engineering; IDSS; MIT

Elchanan Mossel ELMOS@MIT.EDU
Department of Mathematics; IDSS; MIT

M. Amin Rahimian RAHIMIAN@MIT.EDU
IDSS; MIT

Editors: Alina Beygelzimer and Daniel Hsu

Abstract

We study the Bayesian model of opinion exchange of fully rational agents arranged on a network. In this model, the agents receive private signals that are indicative of an unknown state of the world. Then, they repeatedly announce the state of the world they consider most likely to their neighbors, at the same time updating their beliefs based on their neighbors' announcements.

This model is extensively studied in economics since the work of Aumann (1976) and Geanakoplos and Polemarchakis (1982). It is known that the agents eventually agree with high probability on any network. It is often argued that the computations needed by agents in this model are difficult, but prior to our results there was no rigorous work showing this hardness.

We show that it is PSPACE-hard for the agents to compute their actions in this model. Furthermore, we show that it is equally difficult even to approximate an agent's posterior: It is PSPACE-hard to distinguish between the posterior being almost entirely concentrated on one state of the world or another.

Keywords: Bayesian opinion exchange; Computational complexity.

1. Introduction

Background The problem of dynamic opinion exchange is an important field of study in economics, with its roots reaching as far as the Condorcet's jury theorem and, in the Bayesian context, Aumann's agreement theorem. Economists use different opinion exchange models as inspiration for explaining interactions and decisions of market participants. More generally, there is extensive interest in studying how social agents exchange information, form opinions and use them as a basis to make decisions. For a more comprehensive introduction to the subject we refer to surveys addressed to economists (Acemoglu and Ozdaglar, 2011) and mathematicians (Mossel and Tamuz, 2017).

Many models have been proposed and researched, with the properties studied including, among others, if the agents converge to the same opinion, the rate of such convergence, and if the consensus decision is optimal with high probability (this is called *learning*). Two interesting aspects of the differences between models are rules for updating agents' opinions (e.g., fully rational or heuristic) and presence of network structure.

For example, in settings where the updates are assumed to be rational (Bayesian), there is extensive study of models where the agents act in sequence (see, e.g., Banerjee (1992); Bikhchandani et al. (1992); Smith and Sørensen (2000); Acemoglu et al. (2011) for a non-exhaustive selection of works that consider phenomena of *herding* and *information cascades*), as well as models with agents arranged in a network and repeatedly exchanging opinions as time progresses (see some references below). In this work we are interested in the latter class (network models), arguably becoming more relevant given the ubiquity of networks in modern society.

On the other hand, similar questions are studied for models with so-called *bounded rationality*, where the Bayesian updates are replaced with simpler, heuristic rules. Some well-known examples include the DeGroot model (DeGroot, 1974; Golub and Jackson, 2010), the voter model (Clifford and Sudbury, 1973; Holley and Liggett, 1975) and other related variants (Bala and Goyal, 1998; Acemoglu et al., 2010; Arieli et al., 2019).

One commonly accepted reason for studying bounded rationality is that, especially in the network case, Bayesian updates become so complicated as to make fully rational behavior intractable, and therefore unrealistic. However, we are not aware of previous theoretical evidence or formalization of that assertion. Together with another paper of the same authors addressed to economists (Hązła et al., 2018), we consider this work as a development in that direction.

More precisely, we show that computing an agent's opinion in one of the most important and studied Bayesian network models is PSPACE-hard. Furthermore, it is PSPACE-hard even to approximate the rational opinion in any meaningful way. This improves on our NP-hardness result for the same problem shown in Hązła et al. (2018). Our result is tight in the sense that there exists a polynomial space algorithm computing the Bayesian opinion.

Our model and results We are concerned with a certain Bayesian model of opinion exchange and reaching agreement on a network. We are going to call it the *(Bayesian) binary action model*. We consider a network of honest, fully rational agents trying to learn a binary piece of information, e.g., will the price of an asset go up or down, or which political party's policies are more beneficial to the society. We call this information the *state of the world*. Initially, each agent receives an independent piece of information (a *private signal*) that is correlated with the state of the world. According to the principle that "actions speak louder than words", at every time step the agents reveal to their neighbors only one bit: Which of the two possible states they consider more likely. On the other hand, we assume that the agents are honest truth-seekers and always truthfully reveal their preferred state: According to economic terminology they act *myopically* rather than strategically.

More specifically, we assume that the state of the world is encoded in a random variable $\theta \in \{\mathsf{T}, \mathsf{F}\}$ (standing for True and False), distributed according to the uniform prior, shared by all agents. A set of Bayesian agents arranged on a directed graph $G = (V, E)$ performs a sequence of actions at discrete times $t = 0, 1, 2, \ldots$ Before the process starts, each agent u receives a random private signal $\mathcal{S}(u) \in \{0, 1\}$. The collection of random variables $\{\mathcal{S}(u) : u \in V\}$ is independent conditioned on θ. The idea is that $\mathcal{S}(u) = 1$ indicates a piece of evidence for $\theta = \mathsf{T}$ and $\mathcal{S}(u) = 0$ is evidence favoring $\theta = \mathsf{F}$.

At each time $t \geq 0$, the agents simultaneously broadcast *actions* to their neighbors in G. The action $\mathcal{A}(u, t) \in \{\mathsf{T}, \mathsf{F}\}$ is the best guess for the state of the world by agent u at time t: Letting $\mu(u, t)$ be the respective Bayesian posterior probability that $\theta = \mathsf{T}$, the action $\mathcal{A}(u, t) = \mathsf{T}$ if and only if $\mu(u, t) > 1/2$. In subsequent steps, agents update their posteriors based on their neighbors' actions (we assume that everyone is rational, and that this fact and the description of the model are common knowledge) and broadcast updated actions. The process continues indefinitely.

We are interested in computational resources required for the agents to participate in the process described above. That is, we consider the complexity of computing the action $\mathcal{A}(u,t)$ given the private signal $\mathcal{S}(u)$ and history of observations $\{\mathcal{A}(v,t') : v \in \mathcal{N}(u), t' < t\}$, where $\mathcal{N}(u)$ denotes the set of neighbors of u in G. Our main result is that it is worst-case PSPACE-hard for an agent to distinguish between cases where $\mu(u,t) \geq 1 - \exp(-\Theta(N))$ and $\mu(u,t) \leq \exp(-\Theta(N))$, where N is a naturally defined size of the problem. As a consequence, it is PSPACE-hard to compute the action $\mathcal{A}(u,t)$. The result builds on a simpler reduction that shows NP-hardness of the same problem already at time $t = 2$.

Note a *hardness of approximation* aspect of our result: A priori one can imagine a reduction where it is difficult to compute the action $\mathcal{A}(u,t)$ when the Bayesian posterior is close to the threshold $\mu(u,t) \approx 1/2$. However, we demonstrate that it is already hard to distinguish between situations where the posterior is concentrated on one of the extreme values $\mu(u,t) \approx 0$ (and therefore almost certainly $\theta = \mathsf{F}$) and $\mu(u,t) \approx 1$ (and therefore $\theta = \mathsf{T}$).

Our result is tight in the sense that a natural exhaustive search algorithm for computing the action $\mathcal{A}(u,t)$ (see Hązła et al. (2018)) can be implemented in polynomial space. Our hardness results also carry over to other models. In particular, they extend to the case where the signals are continuous, where the prior state of the world is not uniform etc. We also note that we may assume that the agents are never tied or close to tied in their posteriors, see Remark 13.

Related literature Our result improves on our previous work (Hązła et al., 2018), where we showed NP-hardness of approximating Bayesian opinion already at time $t = 2$. In Theorem 1 we describe another NP-hardness reduction. This reduction is different and more complicated than the one in Hązła et al. (2018). However, it is useful in that it allows us to bootstrap an induction leading to a stronger notion of PSPACE-hardness.

One intriguing aspect of our result is a connection to Aumann's agreement theorem. There is a well-known discrepancy (see Cowen and Hanson (2002) for a distinctive take) between reality, where we commonly observe (presumably) honest, well-meaning people "agreeing to disagree", and the Aumann's theorem, stating that this cannot happen for Bayesian agents with common priors and knowledge, i.e., the agents will always end up with the same estimate of the state of the world after exchanging all relevant information. Our result hints at a computational explanation, suggesting that reasonable agreement protocols might be intractable in the presence of network structure. This is notwithstanding some positive computational results of Hanson (2003) and Aaronson (2005), which focus on two agents and come with their own (perhaps unavoidable) caveats.

In particular, Aaronson (2005) considers a setting with two agents, each holding a private signal consisting of n bits. The joint distribution of this $2n$-bit signal is common knowledge of the agents. Their task is to estimate $f(x,y)$ for some function $f : \{0,1\}^{2n} \to [0,1]$. The agents refine their estimates by exchanging messages. Aaronson (2005) studies the protocol where a message consists of the Bayesian expectation $\mathrm{E}[f(x,y)]$ (given everything that the agent knows at the time) with an added noise term. This noise allows for efficient computation in the sense that (assuming that some basic operations on f are efficient) there is an algorithm approximating the estimate $\mathrm{E}[f(x,y)]$ in a number of steps that does not depend on n. Therefore, the protocol exchanging a constant number of messages can be implemented in constant number of steps.

Our setting differs from Aaronson (2005) in that we assume very simple private signals, namely conditionally independent bits. It is clear that in this case a network structure among a large number of agents is necessary for computational hardness. Based on Aaronson's result one could ask if adding noise to agents' messages in our model could make their computations efficient, at least at early time like $t = 2$. It seems to us that restricting to a constant number of agents is essential to Aaronson's analysis and that our proof can be adapted to deal with noise. However, we leave the details to the full version of this paper.

A good deal is known about the model we are considering. From Gale and Kariv (2003) (with an error corrected by Rosenberg et al. (2009), see also similar analysis of earlier, related models in Borkar and Varaiya (1982); Tsitsiklis and Athans (1984)) it follows that if the network G is strongly connected, then the agents eventually converge to the same action (or they become indifferent). The work of Geanakoplos (1994) implies that this agreement is reached in at most $|V| \cdot 2^{|V|}$ time steps. Furthermore, Mossel et al. (2014) showed that in large undirected networks with non-atomic signals, learning occurs: The common agreed action is equal to the state of the world θ, except with probability that goes to zero as the number of agents grows. A good deal remains open, too. For example, it is not known if the $|V| \cdot 2^{|V|}$ bound on the agreement speed can be improved. In this context it is also interesting to note the results of Mossel et al. (2016) who consider a variant of our model with Gaussian structure and revealed beliefs. In contrast to the results presented here, it is shown that in this case, agents' computations are efficient (polynomial time) and convergence time is $O(|V| \cdot \mathrm{diam}(G))$.

We find it interesting that the agents' computations in the binary action model turn out to be not just hard, but PSPACE-hard. PSPACE-hardness of partially observed Markov decision processes (POMDPs) established by Papadimitriou and Tsitsiklis (1987) seems to be a result of a similar kind. On the other hand, there are clear differences: We do not see how to implement our model as POMDP, and embedding a TQBF instance in a POMDP looks more straightforward than what happens in our reduction. Furthermore, and contrary to Papadimitriou and Tsitsiklis (1987), we establish hardness of approximation. We are not aware of many other PSPACE-hardness of approximation proofs. Exceptions are results obtained via PSPACE versions of the PCP theorem (Condon et al., 1995, 1997) and a few other reductions (Marathe et al., 1994; Hunt et al., 1994; Jonsson, 1997, 1999) that concern, among others, some problems on hierarchically generated graphs and an AI-motivated problem of planning in propositional logic.

We note that there are some results on hardness of Bayesian reasoning in static networks in the AI and cognitive science context (see Kwisthout (2018) and its references), but this setting seems quite different from dynamic opinion exchange models. Other related results include work on computational hardness of distributed decision making in control theory (Papadimitriou and Tsitsiklis, 1982; Tsitsiklis and Athans, 1985; Papadimitriou and Tsitsiklis, 1986).

Organization of the paper In Section 2 we give a full description of our model, and state the results precisely. Section 3 contains a discussion of main proof ideas, while in Section 4 there are some suggestions for future work.

The proofs are in the appendices: the NP reduction is described in Section A, the proof of PSPACE-hardness is in Section B, a modification of the proof to use only a fixed number of private signal distributions in Section C, and a proof of #P-hardness in a related *revealed belief* model in Section D.

2. The model and our results

In Section 2.1 we restate the binary action model in more precise terms and introduce some notation. Section 2.2 contains the discussion of our results.

2.1. Binary action model

We consider the *binary action model of Bayesian opinion exchange on a network*. There is a directed graph $G = (V, E)$, the vertices of which we call *agents*. The world has a hidden binary *state* $\theta \in \{\mathsf{T}, \mathsf{F}\}$ with uniform prior distribution. We will analyze a process with discrete time steps $t = 0, 1, 2, \ldots$ At time $t = 0$ each agent u receives a private signal $\mathcal{S}(u) \in \{0, 1\}$. The signals $\mathcal{S}(u)$ are random variables with distributions that are independent across agents after conditioning on θ. Accordingly, the distribution of $\mathcal{S}(u)$ is determined by its *signal probabilities*

$$p_{\theta_0}(u) := \Pr\left[\mathcal{S}(u) = 1 \mid \theta = \theta_0\right], \theta_0 \in \{\mathsf{T}, \mathsf{F}\} .$$

Equivalently, it is determined by its *log-likelihoods ratios (LLRs)*

$$\ell_b(u) = \ln \frac{\Pr[\mathcal{S}(u) = b \mid \theta = \mathsf{T}]}{\Pr[\mathcal{S}(u) = b \mid \theta = \mathsf{F}]} = \ln \frac{\Pr[\theta = \mathsf{T} \mid \mathcal{S}(u) = b]}{\Pr[\theta = \mathsf{F} \mid \mathcal{S}(u) = b]}, b \in \{0, 1\} .$$

Note that there is a one-to-one correspondence between probabilities p_T and p_F with $p_\mathsf{T} \neq p_\mathsf{F}$, and LLRs ℓ_0, ℓ_1 with $\ell_0 \cdot \ell_1 < 0$. We will always assume that a signal $\mathcal{S}(u) = 1$ is evidence towards $\theta = \mathsf{T}$ and vice versa. This is equivalent to saying that $p_\mathsf{T} > p_\mathsf{F}$ or $\ell_1 > 0$ and $\ell_0 < 0$. We allow some agents to not receive private signals: This can be "simulated" by giving them non-informative signals with $p_\mathsf{T}(u) = p_\mathsf{F}(u)$. We will refer to all signal probabilities taken together as the *signal structure* of the Bayesian network. A specific pattern of signals $s \in \{0, 1\}^{|V'|}$ (where V' denotes the subset of agents that receive informative signals) will be called a *signal configuration*.

We assume that all this structural information is publicly known, but the agents do not have direct access to θ or others' private signals. Agents are presumed to be rational, to know that everyone else is rational, to know that everyone knows, etc. (common knowledge of rationality). At each time $t \geq 0$, we define $\mu(u, t)$ to be the *belief* of agent u: The conditional probability that $\theta = \mathsf{T}$ given everything that u observed at times $t' < t$. More precisely, letting $\mathcal{N}(u)$ be the (out)neighbors of u in G and defining

$$H(u, t) := \left\{\mathcal{A}(v, t') : t' < t, v \in \mathcal{N}(u)\right\} .$$

as the *observation history* of agent u we let $\mu(u, t) := \Pr[\theta = \mathsf{T} \mid \mathcal{S}(u), H(u, t)]$. Accordingly, if $(u, v) \in E(G)$ we will say that agent u *observes* agent v.

Agent u broadcasts to its in-neighbors the *action* $\mathcal{A}(u, t) \in \{\mathsf{T}, \mathsf{F}\}$, which is the state of the world that u considers more likely according to $\mu(u, t)$ (assume that ties are broken in an arbitrary deterministic manner, say, in favor of F). Then, the protocol proceeds to time step $t + 1$ and the agents update their beliefs and broadcast updated actions. The process continues indefinitely. Note that the beliefs and actions become deterministic once the private signals are fixed.

The first two time steps of the process are relatively easy to understand: At time $t = 0$ an agent broadcasts $\mathcal{A}(u, 0) = \mathsf{T}$ if and only if $\mathcal{S}(u) = 1$ and the belief $\mu(u, 0)$ can be easily computed from the LLR $\ell_{\mathcal{S}(u)}(u)$. At time $t = 1$, an agent broadcasts $\mathcal{A}(u, 1) = \mathsf{T}$ if and only if

$$\ell_{\mathcal{S}(u)}(u) + \sum_{v \in \mathcal{N}(u)} \ell_{\mathcal{S}(v)}(v) > 0 , \tag{1}$$

where the private signals $\mathcal{S}(v)$ can be inferred from observed actions $\mathcal{A}(v, 0)$. The sum (1) determines the LLR associated with belief $\mu(u, 1)$. However, at later times the actions of different neighbors are not independent anymore and accounting for those dependencies seems difficult.

Let Π be a Bayesian network, i.e., a directed graph $G = (V, E)$ together with the signal structure. We do not commit to any particular representation of probabilities of private signals. Our reduction remains valid for any reasonable choice. We are interested in the hardness of computing the actions that the agents need to broadcast. More precisely, we consider the complexity of computing the function

$$\text{BINARY-ACTION}(\Pi, t, u, \mathcal{S}(u), H(u,t)) := \mathcal{A}(u, t)$$

that computes the action $\mathcal{A}(u, t)$ given the Bayesian network Π, time t, agent u, its private signal $\mathcal{S}(u)$ and observation history $H(u, t)$. Relatedly, we will consider computing the belief

$$\text{BINARY-BELIEF}(\Pi, t, u, \mathcal{S}(u), H(u,t)) := \mu(u, t) \ .$$

Note that computing BINARY-ACTION is equivalent to distinguishing between BINARY-BELIEF $>$ $1/2$ and BINARY-BELIEF $\leq 1/2$.

2.2. Our results

Our first result implies that computing BINARY-ACTION at time $t = 2$ is NP-hard. We present it as a standalone theorem, since the NP reduction and its analysis are used as a building block in the more complicated PSPACE reduction.

Theorem 1 *There exists an efficient reduction from a* 3-SAT *formula* ϕ *with* N *variables and* M *clauses to an input of* BINARY-ACTION$(\Pi, t, u, H(u,t))$ *such that:*

- *The size (number of agents and edges) of the Bayesian graph* G *is* $O(N + M)$*, the time is set to* $t = 2$ *and agent* u *does not receive a private signal.*
- *All probabilities of private signals are efficiently computable real numbers satisfying*

$$\exp(-(O(N)) \leq p_{\theta_0}(v) \leq 1 - \exp(-O(N)),\ v \in V,\ \theta_0 \in \{\mathsf{T}, \mathsf{F}\} \ .$$

- *If* ϕ *is satisfiable, then the belief* $\mu(u, 2)$ *satisfies*

$$\mu(u, 2) = 1 - \exp(-\Theta(N)) \ .$$

- *If* ϕ *is not satisfiable, then we have*

$$\mu(u, 2) = \exp(-\Theta(N)) \ .$$

Corollary 2 *Both distinguishing between the cases* BINARY-BELIEF $> 1 - \exp(-\Omega(N))$ *and* BINARY-BELIEF $< \exp(-\Omega(N))$ *and computing* BINARY-ACTION *are* NP*-hard. Furthermore, since meanings of the labels* T *and* F *can be reversed, computing* BINARY-ACTION *is also* coNP*-hard.*

Our main result improves Theorem 1 to PSPACE-hardness. It is a direct reduction from the canonical PSPACE-complete language of true quantified Boolean formulas TQBF.

Theorem 3 *There exists an efficient reduction from a* TQBF *formula* Φ

$$\Phi = Q_K \boldsymbol{x}_K \cdots \exists \boldsymbol{x}_1 \phi(\boldsymbol{x}_K, \dots, \boldsymbol{x}_1) ,$$

where ϕ is a 3-CNF formula with N variables and M clauses, there are K variable blocks with alternating quantifiers and the last quantifier is existential, to an input of the computational problem $\text{BINARY-ACTION}(\Pi, t, u, H(u,t))$ *such that:*

- *The number of agents in the Bayesian graph G is $O(N^2(N+M))$, the time is set to $t = 2K$ and agent u does not receive a private signal.*
- *All probabilities of private signals are efficiently computable real numbers satisfying*

$$\exp(-(O(N)) \le p_{\theta_0}(v) \le 1 - \exp(-O(N)),\ v \in V,\ \theta_0 \in \{\mathsf{T}, \mathsf{F}\} . \tag{2}$$

- *If Φ is true, then $\mu(u, 2K) = 1 - \exp(-\Theta(N))$. If Φ is false, then $\mu(u, 2K) = \exp(-\Theta(N))$.*

Corollary 4 *Both distinguishing between the cases* $\text{BINARY-BELIEF} > 1 - \exp(-\Omega(N))$ *and* $\text{BINARY-BELIEF} < \exp(-\Omega(N))$ *and computing* BINARY-ACTION *are* PSPACE*-hard.*

Remark 5 *Note that the statement of Theorem 3 immediately gives Σ_K-hardness of approximating* BINARY-BELIEF *at time $t = 2K$. Furthermore, due to the symmetric nature of the problem (cf. Corollary 2), approximating* BINARY-BELIEF *at time $t = 2K$ is also Π_K-hard.*

Remark 6 *For ease of exposition we define networks in the reductions to be directed, but due to additional structure that we impose (see paragraph "Network structure and significant times" in Section A) it is easy to see that they can be assumed to be undirected. This is relevant insofar as a strong form of learning occurs only on undirected graphs (see Mossel et al. (2014)).*

We briefly discuss the upper bounds for computing BINARY-BELIEF. As we explained in Section 2.1, for $t = 0, 1$ there is a simple formula to compute an agent's belief. If we assume $t = \text{poly}(n)$, then a natural exhaustive search algorithm (see Hązła et al. (2018)) can be implemented in polynomial space, matching the lower bound from Theorem 3. On the other hand, we do not know how to improve on this algorithm even for $t = 2$. We note, however, that since Corollary 2 establishes both NP-hardness and coNP-hardness, the problem for $t = 2$ is unlikely to be NP-complete.

One possible objection to Theorem 3 is that it uses signal distributions with probabilities exponentially close to zero and one. We do not think this is a significant issue, and it helps avoid some technicalities. Nevertheless, in Section C we prove a version of Theorem 3 where all private signals come from a fixed family of, say, at most fifty distributions. This is at the cost of a (non-asymptotic) increase in the size of the graph.

Theorem 7 *The reduction from Theorem 3 can be modified such that all private signals come from a fixed family of at most fifty distributions.*

Remark 8 *It is possible to modify our proofs to give hardness of distinguishing between $\mu(u,t) \le \exp(-\Omega(N^K))$ and $\mu(u,t) \ge 1 - \exp(-\Omega(N^K))$ for any constant K (recall that N is the number of variables in the formula ϕ). This is at the cost of allowing signal probabilities in the range*

$$\exp(-(O(N^K)) \le p_{\theta_0}(v) \le 1 - \exp(-O(N^K))$$

or, in the bounded signal case, increasing the network size to $O(N^{K+2})$. Consequently, in the latter case we get hardness of approximation up to an $\exp(O(|V|^\alpha))$ factor for any constant $\alpha < 1$, where $|V|$ is the number of agents.

3. Main proof ideas

NP**-hardness** The NP-hardness proof in Section A consists of an analysis of a composition of several gadgets. We will think of the agent u from input to BINARY-ACTION as "observer" and accordingly call it OBS. The Bayesian network features gadgets that represent variables and clauses. The private signals in variable gadgets correspond to assignments $\boldsymbol{w}$ to the formula ϕ. Furthermore, there is an "evaluation agent" EVAL that interacts with all clause gadgets. We use more gadgets that "implement" counting to ensure that what OBS observes is consistent with one of two possible kinds of signal configurations:

- $\mathcal{S}(\mathsf{EVAL}) = 1$ and the signals of variable agents correspond to an arbitrary assignment $\boldsymbol{w}$.
- $\mathcal{S}(\mathsf{EVAL}) = 0$ and the signals of variable agents correspond to a *satisfying* assignment $\boldsymbol{w}$.

Then, we use another gadget to "amplify" the information that is conveyed about the state of the world by the signal $\mathcal{S}(\mathsf{EVAL})$. If ϕ has no satisfying assignment, then $\mathcal{S}(\mathsf{EVAL}) = 1$ and this becomes amplified to a near-certainty that $\theta = \mathsf{F}$ (for technical reasons this is the opposite conclusion than suggested by $\mathcal{S}(\mathsf{EVAL}) = 1$). On the other hand, we design the signal structure such that even a single satisfying assignment tips the scales and amplifies to $\theta = \mathsf{T}$ with high probability (whp).

We note that one technical challenge in executing this plan is that some of our gadgets are designed to "measure" (e.g., count) certain properties of the private signals, but these measurements use auxiliary agents with their own signals, affecting Bayesian beliefs. We need to be careful to cancel out these unintended effects at every step.

PSPACE**-hardness** The high-level idea to improve on the NP-hardness proof is that once we know that agents can solve hard problems, we can use them to help the observer agent solve an even harder problem. Of course this has to be done in a careful way, since the answer to a partial problem cannot be directly revealed to the observer (the whole point is that we do not know a priori what this answer is).

We will show PSPACE-hardness by reduction from the canonical PSPACE-complete language TQBF. More precisely, we use a representation of quantified Boolean formulas

$$\Phi = Q_K \boldsymbol{x}_K \cdots Q_1 \boldsymbol{x}_1 : \phi(\boldsymbol{x}_K, \ldots, \boldsymbol{x}_1) ,$$

where:

- Q_i is a quantifier such that $Q_i \in \{\exists, \forall\}$, $Q_i \neq Q_{i+1}$ and $Q_1 = \exists$.
- $\boldsymbol{x}_K, \ldots, \boldsymbol{x}_1$ are *blocks of variables* such that their total count is $|\boldsymbol{x}_K| + \ldots + |\boldsymbol{x}_1| = N$.
- ϕ is a 3-CNF formula over variables $\boldsymbol{x}_K, \ldots, \boldsymbol{x}_1$ with M clauses.

The language TQBF consists of all formulas Φ that are true. It is common and useful to think of Φ as defining a "position" in a game, where "Player 1" chooses values of variables under existential quantifiers, "Player 0" chooses values of variables under universal quantifiers, and the objective of Player s is to evaluate ϕ to the value s. Under that interpretation, $\Phi \in \mathrm{TQBF}$ if and only if Player 1 has a winning strategy in the given position.

Keeping that in mind, we can give an intuition for the proof: In the 3-SAT reduction, if the formula had a satisfying assignment, then agent OBS could conclude whp. that the "hidden" assignment is satisfying, and $\theta = \mathsf{T}$. Otherwise, the hidden assignment is not satisfying and $\theta = \mathsf{F}$ whp. In the PSPACE reduction, the hidden assignment will correspond (whp.) to a "transcript" of the game played according to a winning strategy for one of the players, and θ will be determined by the winning player. This will be achieved by implementing a sequence of observer agents $\mathsf{OBS}_1, \ldots, \mathsf{OBS}_K$, where:

- Ultimately, the hardness will be shown for the computation of agent OBS_K.
- Agent OBS_i directly observes variable agents in blocks $\boldsymbol{x}_K, \ldots, \boldsymbol{x}_{i+1}$. It will be useful to think of OBS_i as "computing" a logical formula with $i-1$ quantifier switches (since from the perspective of this agent blocks $\boldsymbol{x}_K, \ldots, \boldsymbol{x}_{i+1}$ are set to fixed values).
- For each i, there is a (slightly more complicated) gadget similar to the "EVAL-gadget" employed in the 3-SAT reduction. This gadget involves OBS_{i-1} as well as two new agents B_i and C_i and is observed by OBS_i. Its purpose is to "flip" relative LLRs of different types of variable assignments to implement a quantifier switch.

4. Conclusion

A natural open question is to make progress on the approximate-case hardness in one of the models. For example, one could try to establish NP-hardness for a worst-case network, but holding for signal configurations arising with non-negligible probability. This might require significant new ideas.

Another interesting problem arises from trying to extend our results to the revealed belief model, as discussed in Section D. Thinking in terms of games, consider a class of "no-mistakes-allowed" games: Games where the player with winning strategy always has exactly one winning move, with all alternative moves in a given position leading to a losing position (and this property holding recursively in all positions reachable from the initial one).

Certainly deciding if a position is winning for the first player in such games is in PSPACE. On the other hand, since such a game with all moves performed by the existential player corresponds to a SAT formula with zero or one satisfying assignments, by the Valiant-Vazirani theorem (Valiant and Vazirani, 1986) it is also (morally) NP-hard. This leaves a large gap between NP and PSPACE.

For example, suppose we want to prove Π_2-hardness in the revealed belief model. Then it is natural to consider formulas of the form $\forall \boldsymbol{x} \exists \boldsymbol{y} : \phi(\boldsymbol{x}, \boldsymbol{y})$, and the question becomes: How hard is it to distinguish between the cases:

- YES: For all $\boldsymbol{x}$, there exists unique $\boldsymbol{y}_{\boldsymbol{x}}$ such that $\phi(\boldsymbol{x}, \boldsymbol{y}_{\boldsymbol{x}}) = 1$.
- NO: There exists unique $\boldsymbol{x}_0$ such that for all $\boldsymbol{y}$ we have $\phi(\boldsymbol{x}_0, \boldsymbol{y}) = 0$. For all other $\boldsymbol{x}$, there exists unique $\boldsymbol{y}_{\boldsymbol{x}}$ such that $\phi(\boldsymbol{x}, \boldsymbol{y}_{\boldsymbol{x}}) = 1$.

How hard is this problem? In particular, can it be shown to be harder than NP (in some sense)? Hardness of such games can be thought of as a generalization of the Valiant-Vazirani theorem.

Acknowledgments

We thank Ehsan Emamjomeh-Zadeh for pointing out the PSPACE upper bound. We are also grateful to referees for helpful corrections. This work was partially supported by awards ONR N00014-16-1-2227, NSF CCF-1665252 and ARO MURI W911NF-12-1-0509 and by a Vannevar Bush Fellowship.

References

Scott Aaronson. The complexity of agreement. In *ACM Symposium on Theory of Computing*, pages 634–643, 2005. doi: 10.1145/1060590.1060686.

Daron Acemoglu and Asuman Ozdaglar. Opinion dynamics and learning in social networks. *Dynamic Games and Applications*, 1(1):3–49, 2011. doi: 10.1007/s13235-010-0004-1.

Daron Acemoglu, Asuman Ozdaglar, and Ali ParandehGheibi. Spread of (mis)information in social networks. *Games and Economic Behavior*, 70(2):194–227, 2010. doi: 10.1016/j.geb.2010.01.005.

Daron Acemoglu, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar. Bayesian learning in social networks. *The Review of Economic Studies*, 78(4):1201–1236, 2011. doi: 10.1093/restud/rdr004.

Itai Arieli, Yakov Babichenko, and Manuel Mueller-Frank. Naive learning through probability matching. SSRN preprint, doi: 10.2139/ssrn.3338015, 2019.

Robert J. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976. doi: 10.1214/aos/1176343654.

Venkatesh Bala and Sanjeev Goyal. Learning from neighbours. *The Review of Economic Studies*, 65(3):595–621, 1998. doi: 10.1111/1467-937X.00059.

Abhijit V. Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, 107 (3):797–817, 1992. doi: 10.2307/2118364.

Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026, 1992. doi: 10.1086/261849.

Vivek Borkar and Pravin P. Varaiya. Asymptotic agreement in distributed estimation. *IEEE Transactions on Automatic Control*, 27(3):650–655, 1982. doi: 10.1109/TAC.1982.1102982.

Peter Clifford and Aidan Sudbury. A model for spatial conflict. *Biometrika*, 60(3):581–588, 1973. doi: 10.1093/biomet/60.3.581.

Anne Condon, Joan Feigenbaum, Carsten Lund, and Peter Shor. Probabilistically checkable proof systems and nonapproximability of *PSPACE*-hard functions. *Chicago Journal of Theoretical Computer Science*, 1995(4), 1995.

Anne Condon, Joan Feigenbaum, Carsten Lund, and Peter Shor. Random debaters and the hardness of approximating stochastic functions. *SIAM Journal on Computing*, 26(2):369–400, 1997. doi: 10.1137/S0097539793260738.

Tyler Cowen and Robin Hanson. Are disagreements honest? Technical report, George Mason University Economics, 2002.

Morris H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69 (345):118–121, 1974. doi: 10.1080/01621459.1974.10480137.

Douglas Gale and Shachar Kariv. Bayesian learning in social networks. *Games and Economic Behavior*, 45(2):329–346, 2003. doi: https://doi.org/10.1016/S0899-8256(03)00144-1.

John Geanakoplos. Common knowledge. In *Handbook of Game Theory with Economic Applications*, volume 2, chapter 40, pages 1437–1496. Elsevier, 1st edition, 1994. doi: 10.1016/S1574-0005(05)80072-4.

John D. Geanakoplos and Heraklis M. Polemarchakis. We can't disagree forever. *Journal of Economic Theory*, 28(1):192–200, 1982. doi: 10.1016/0022-0531(82)90099-0.

Benjamin Golub and Matthew O. Jackson. Naïve learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–149, February 2010. doi: 10.1257/mic.2.1.112.

Robin Hanson. For Bayesian wannabes, are disagreements not about information? *Theory and Decision*, 54(2):105–123, 2003. doi: 10.1023/A:1026251404228.

Richard A. Holley and Thomas M. Liggett. Ergodic theorems for weakly interacting infinite systems and the voter model. *The Annals of Probability*, 3(4):643–663, 1975. doi: 10.1214/aop/1176996306.

Harry B. Hunt, Madhav V. Marathe, and Richard E. Stearns. Generalized CNF satisfiability problems and non-efficient approximability. In *Complexity Theory Conference (CCC)*, pages 356–366, 1994.

Jan Hązła, Ali Jadbabaie, Elchanan Mossel, and M. Amin Rahimian. Bayesian decision making in groups is hard. arXiv:1705.04770v3, 2018.

Peter Jonsson. A nonapproximability result for finite function generation. *Information Processing Letters*, 63(3):143–145, 1997. doi: 10.1016/S0020-0190(97)00112-9.

Peter Jonsson. Strong bounds on the approximability of two PSPACE-hard problems in propositional planning. *Annals of Mathematics and Artificial Intelligence*, 26(1):133–147, 1999. doi: 10.1023/A:1018954827926.

Johan Kwisthout. Approximate inference in Bayesian networks: Parameterized complexity results. *International Journal of Approximate Reasoning*, 93:119–131, 2018. doi: https://doi.org/10.1016/j.ijar.2017.10.029.

Madhav V. Marathe, Harry B. Hunt, and S. S. Ravi. The complexity of approximating PSPACE-complete problems for hierarchical specifications. *Nordic Journal of Computing*, 1(3):275–316, 1994.

Elchanan Mossel and Omer Tamuz. Opinion exchange dynamics. *Probability Surveys*, 14:155–204, 2017. doi: 10.1214/14-PS230.

Elchanan Mossel, Allan Sly, and Omer Tamuz. Asymptotic learning on Bayesian social networks. *Probability Theory and Related Fields*, 158(1):127–157, 2014. doi: 10.1007/s00440-013-0479-y.

Elchanan Mossel, Noah Olsman, and Omer Tamuz. Efficient Bayesian learning in social networks with Gaussian estimators. In *Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 425–432, 2016. doi: 10.1109/ALLERTON.2016.7852262.

Christos H. Papadimitriou and John Tsitsiklis. On the complexity of designing distributed protocols. *Information and Control*, 53(3):211–218, 1982. doi: 10.1016/S0019-9958(82)91034-8.

Christos H. Papadimitriou and John Tsitsiklis. Intractable problems in control theory. *SIAM Journal on Control and Optimization*, 24(4):639–654, 1986. doi: 10.1137/0324038.

Christos H. Papadimitriou and John N. Tsitsiklis. The complexity of Markov decision processes. *Mathematics of Operations Research*, 12(3):441–450, 1987. doi: 10.1287/moor.12.3.441.

Dinah Rosenberg, Eilon Solan, and Nicolas Vieille. Informational externalities and emergence of consensus. *Games and Economic Behavior*, 66(2):979–994, 2009. doi: 10.1016/j.geb.2008.09.027.

Lones Smith and Peter Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000. doi: 10.1111/1468-0262.00113.

John N. Tsitsiklis and Michael Athans. Convergence and asymptotic agreement in distributed decision problems. *IEEE Transactions on Automatic Control*, 29(1):42–50, 1984. doi: 10.1109/TAC.1984.1103385.

John N. Tsitsiklis and Michael Athans. On the complexity of decentralized decision making and detection problems. *IEEE Transactions on Automatic Control*, 30(5):440–446, 1985. doi: 10.1109/TAC.1985.1103988.

Leslie G. Valiant and Vijay V. Vazirani. NP is as easy as detecting unique solutions. *Theoretical Computer Science*, 47:85–93, 1986. doi: 10.1016/0304-3975(86)90135-0.

Appendix A. NP-hardness: Proof of Theorem 1

In this section we describe and analyze the reduction from 3-SAT to BINARY-ACTION, proving Theorem 1. The reduction is used as a building block in the PSPACE-hardness proof, but it is also useful in terms of developing intuition for the more technical proof of Theorem 3. We proceed by explaining gadgets that we use, describing how to put them together in the full reduction, and finally proving its correctness.

Threshold gadget Say there are agents $v_1, \ldots, v_K$ that do not observe anyone and receive private signals $\mathcal{S}(v_i)$ with respective LLRs $\ell_0(v_i)$ and $\ell_1(v_i)$. Additionally, there is an observer agent OBS and we would like to reveal to it, at time $t = 1$, that the sum of LLRs of agents $v_1, \ldots, v_K$ exceeds some threshold δ:

$$L := \sum_{i=1}^{K} \ell_{\mathcal{S}(v_i)}(v_i) > \delta \; ,$$

without disclosing anything else about the private signals.[1] This is achieved by the gadget in Figure 1.

Figure 1: Threshold gadget.

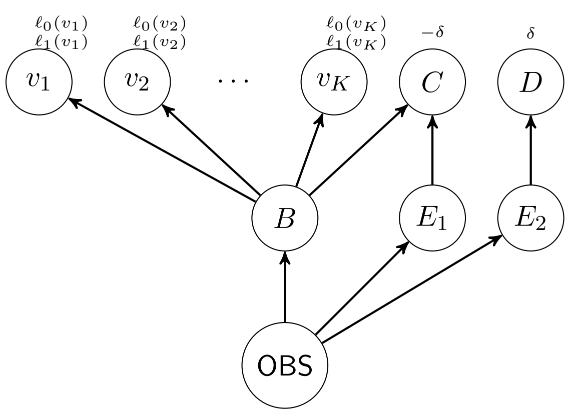

Figure 2: Notation for the threshold gadget.

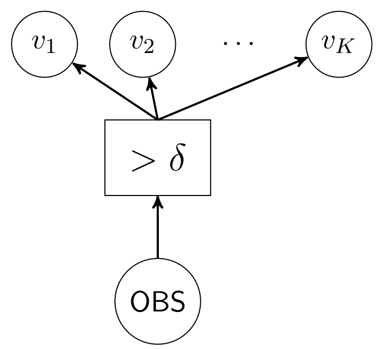

We describe the gadget for $\delta > 0$. Agent C receives a private signal with $\ell_0(C) = -\delta$ (and arbitrary $\ell_1(C)$) and agent D with $\ell_1(D) = \delta$. Agents B, E_1 and E_2 (we will call the latter two the "dummy" agents) do not receive private signals.

We stress that our overall reduction will demonstrate the hardness of computation for agent OBS. Therefore, we need to specify the observation history of OBS. By our tie-breaking convention, it must be $\mathcal{A}(B, 0) = \mathcal{A}(E_1, 0) = \mathcal{A}(E_2, 0) = \mathsf{F}$. Furthermore, we specify $\mathcal{A}(B, 1) = \mathcal{A}(E_2, 1) = \mathsf{T}$ and $\mathcal{A}(E_1, 1) = \mathsf{F}$.

Based on that information, agent OBS can infer that $\mathcal{S}(C) = 0$, $\mathcal{S}(D) = 1$ and, since the action $\mathcal{A}(B, 2)$ is determined by the sign of $L - \delta$, that $L > \delta$. The purpose of agent D is to counteract the

1. We assume that δ is chosen such that $L = \delta$ never happens.

effect of this "measurement" on the estimate of the state of the world by OBS. More precisely, let

$$P(s_1, \ldots, s_K, \theta_0) := \Pr\left[\bigwedge_{i=1}^{K} \mathcal{S}(v_i) = s_i \wedge \theta = \theta_0\right] , \tag{3}$$

$$P(s_1, \ldots, s_K, s_C, s_D, \theta_0) := \Pr\left[\bigwedge_{i=1}^{K} \mathcal{S}(v_i) = s_i \wedge \mathcal{S}(C) = s_C \wedge \mathcal{S}(D) = s_D \wedge \theta = \theta_0\right] . \tag{4}$$

Based on the discussion above, we have the following:

Claim 9 *Let $s_1, \ldots, s_K$ be private signals of $v_1, \ldots, v_K$. Similarly, let (s_C, s_D) be private signals of C and D. Then:*

- *If $\sum_{i=1}^{K} \ell_{s_i}(v_i) < \delta$, then there are no signals (s_C, s_D) that make $(s_1, \ldots, s_K, s_C, s_D)$ consistent with observations of OBS.*
- *If $\sum_{i=1}^{K} \ell_{s_i}(v_i) > \delta$, then there exists unique configuration (s_C, s_D) consistent with observations of OBS and the (prior) probability of this configuration when the state is θ_0 is*

$$P(s_1, \ldots, s_K, s_C, s_D, \theta_0) = P(s_1, \ldots, s_K, \theta_0) \cdot \alpha , \tag{5}$$

where $\alpha := (1-p_{\mathsf{T}}(C))p_{\mathsf{T}}(D) = e^{\ell_0(C)+\ell_1(D)}(1-p_{\mathsf{F}}(C))p_{\mathsf{F}}(D) = e^{-\delta+\delta}(1-p_{\mathsf{F}}(C))p_{\mathsf{F}}(D) = (1-p_{\mathsf{F}}(C))p_{\mathsf{F}}(D)$ does not depend on θ_0.

Similar reasoning can be used for the case when $\delta < 0$ and/or checking the opposite inequality $L < \delta$. We will say that an agent OBS observes a threshold gadget if it observes agents B, E_1 and E_2 and denote it as shown in Figure 2. Note that in our diagrams we use circles to denote agents and boxes to denote gadgets. The latter typically contain several auxiliary agents.

Network structure and significant times It might appear that the threshold gadget is more complicated than needed. The reason for this is that we will impose certain additional structure on the graph to facilitate a more detailed analysis, which will be later used in the proof of Theorem 3. Specifically, we will always make sure that the graph is a DAG, with only the observer agent having in-degree zero. Furthermore, an agent will receive a private signal if and only if its out-degree is zero (recall a directed edge $B \to C$ indicates that B observes C).

Furthermore, we will arrange the graph such that each agent will learn new information at a single, fixed time step. That is, for every agent B there will exist a *significant time* $t(B)$ such that $\mu(B, t') = 1/2$ for $t' < t(B)$ and $\mu(B, t') = \mu(B, t(B))$ for $t' > t(B)$. If B receives a private signal, then $t(B) = 0$. Otherwise, $t(B)$ is determined by the (unique) path length from B to an out-degree zero agent. For example, in Figure 1 significant times are $t(B) = t(E_1) = t(E_2) = 1$ and $t(\mathsf{OBS}) = 2$.

Accordingly, we will use notation $\mu(B)$ and $\mathcal{A}(B)$ to denote agent beliefs and actions at the significant time. Let B and C be agents with $t(B) < t(C) - 1$. In the following, we will sometimes say that C observes B, even though that would contradict the significant time requirement (a direct edge $C \to B$ implies that $t(B) = t(C) - 1$). Whenever we do so, it should be understood that there is a path of "dummy" nodes of appropriate length between B and C (cf. E_1 and E_2 in Figure 1). For clarity, we will omit dummy nodes from the figures.

Counting gadget Assume now that the agents $v_1, \dots, v_K$ receive private signals with identical LLRs $\ell_0 < 0$ and $\ell_1 > 0$ and that a number k, $1 \le k \le K$ is given. Then, building on the threshold gadget, we can convey the information that exactly k out of K agents received private signal 1. Letting $\delta := K\ell_0 + (k - 0.5)(\ell_1 - \ell_0)$ and $\delta' := \delta + \ell_1 - \ell_0$, we compose two threshold gadgets as shown in Figure 3.

Figure 3: Counting gadget.

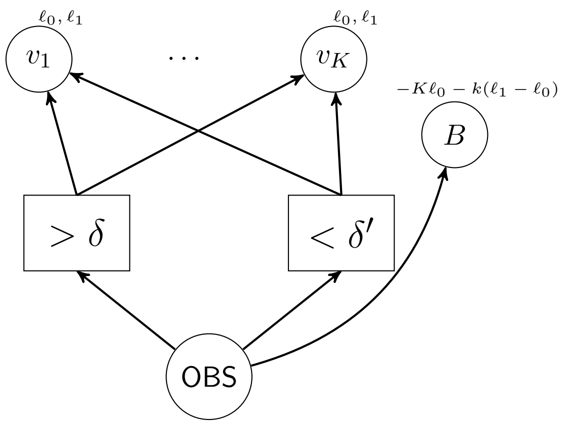

Figure 4: Two counting gadgets illustrating the notation. The equivalence symbol on the right-hand gadget denotes presence of the optional agent B.

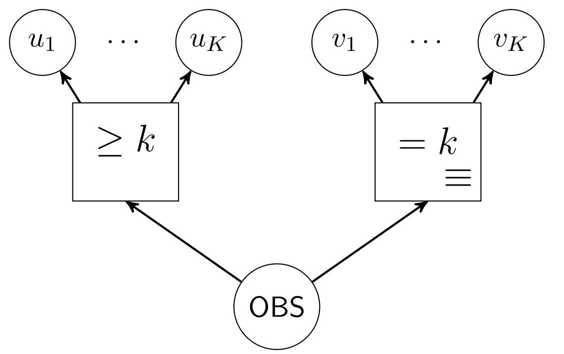

Agent B shown in Figure 3 is optional: Depending on our needs we will use the counting gadget with or without it. It is used to preserve the original belief of OBS after learning the count of private signals of agents v_i. It receives a private signal with $\ell_b(B) := \ell := -K\ell_0 - k(\ell_1 - \ell_0)$ for appropriate $b \in \{0, 1\}$ (depending on the sign of ℓ) and broadcasts the corresponding value of θ_0.

Again, we emphasize that the gadget will be used in the context of the reduction, where we specify the observation history of agent OBS. Therefore, the definition of the gadget includes the values of private signals of agent B, as well as auxiliary agents in the counting gadgets, but not the values of private signals of agents $v_1, \dots, v_K$.

By similar analysis as for the threshold gadget and using the $P(\cdot)$ notation as in (3)–(5) we we can establish a formal claim similar to Claim 9:

Claim 10 *Consider a counting gadget as shown in Figure 3. Let $s_1, \dots, s_K$ be private signals of agents $v_1, \dots, v_K$. Let $\boldsymbol{s}$ represent private signals of all auxiliary agents in the threshold gadgets and s_B a private signal of agent B.*

Then, the only configurations $s_1, \dots, s_K$ consistent with observations of OBS are those for which $\sum_{i=1}^{K} s_i = k$. Furthermore, for any such configuration there exists a unique configuration $\boldsymbol{s}$ (and s_B, if agent B is present) such that (depending on the presence of B):

$$P(s_1, \dots, s_K, \boldsymbol{s}, \theta_0) = P(s_1, \dots, s_K, \theta_0) \cdot \alpha = P(\theta_0) \cdot \alpha \;, \tag{6}$$

$$P(s_1, \dots, s_K, \boldsymbol{s}, s_B, \theta_0) = \beta \;, \tag{7}$$

where $\alpha := \alpha(k, K, \ell_0, \ell_1) > 0$ is easily computable and does not depend on $s_1, \dots, s_K$ or θ_0, but the value of the other term $P(\theta_0)$ is in general dependent on θ_0. On the other hand, if B is present, then $\beta := \beta(k, K, \ell_0, \ell_1) > 0$ does not depend at all on the private signals or state of the world.

Remark 11 *The same technique can be used to obtain inequalities (e.g., checking that at least k out of K private signals are ones). However, note that in case of inequality only the version without agent B, achieving guarantee (6), can be implemented. This is because in order to get (7) we use the knowledge of the exact LLR shift induced by the fact that exactly k agents received positive signals, i.e., $K\ell_0 + k(\ell_1 - \ell_0)$. Of course we lack this information if we only know that* at least *k agents received positive signals.*

We will say that an agent OBS observes the counting gadget if it observes both respective threshold gadgets (and B, if present). We will denote counting gadgets as in Figure 4.

Not-equal gadget Another related gadget that we will use reveals to the observer that two agents u, v with LLRs ℓ_0, ℓ_1 and m_0, m_1, respectively, receive opposite signals $\mathcal{S}(u) \neq \mathcal{S}(v)$. Since $\ell_0, m_0 < 0 < \ell_1, m_1$, this is achieved by using two threshold gadgets to check that $\ell_0 + m_0 < \ell_{\mathcal{S}(u)} + m_{\mathcal{S}(v)} < \ell_1 + m_1$, where we set the thresholds in the threshold gadgets as $\ell_0 + m_0 + \varepsilon$ and $\ell_1 + m_1 - \varepsilon$ for an appropriately small $\varepsilon > 0$. We will denote the not-equal gadget as in Figure 5.

Figure 5: Not-equal gadget.

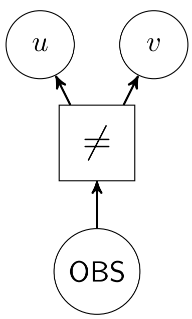

Figure 6: Variable gadget.

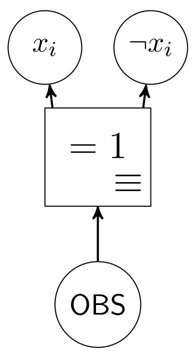

Figure 7: Clause gadget.

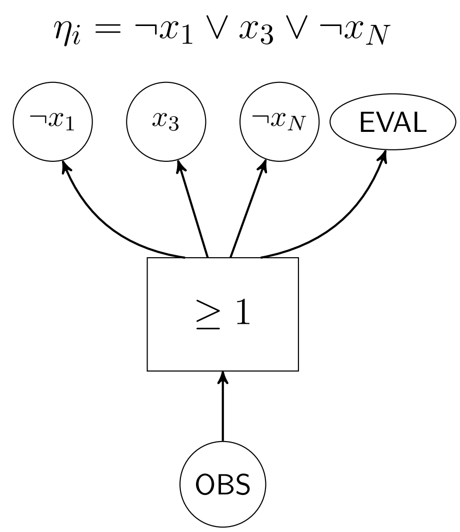

Variable and clause gadgets Our reduction is from the standard form of 3-SAT, where we are given a CNF formula on N variables $x_1, \ldots, x_N$. The formula is a conjunction of M clauses $\eta_1, \ldots, \eta_M$, where each clause is a disjunction of exactly three literals on distinct variables.

We introduce two global agents. One of them is called OBS and we mean it as an "observer agent". This is the agent for which we establish hardness of computation. We will follow the rule that OBS observes all gadgets that are present in the network. Second, we place an "evaluation agent" EVAL with private signal probabilities $p_\mathsf{T} := 0.9$ and $p_\mathsf{F} := 0.4$.

Furthermore, for each variable in the CNF formula, we introduce two agents x_i and $\neg x_i$ that receive private signals given by p_T and p_F. Then, we encompass those two agents in a counting gadget with equivalence as shown in Figure 6.

For each clause η_i, we introduce a counting gadget on four agents: Three agents corresponding to the literals in the clause (note that they are observed directly and not through the variable gadgets), and the EVAL agent. The gadget ensures that at least one of those agents received signal 1. An illustration is provided in Figure 7.

The reduction We put the agents EVAL and OBS and the variable and clause gadgets together, as explained in previous paragraphs. Finally, we add two more agents B and C. We will choose a natural number $r := \gamma \cdot N$ for an absolute big enough constant $\gamma > 0$. Agent B receives private signals with $p_{\mathsf{T}}(B) = 1 - \alpha_1^r$ and $p_{\mathsf{F}}(B) = \alpha_2^r$ and agent C with $p_{\mathsf{T}}(C) = 1 - \alpha_3^r$ and $p_{\mathsf{F}}(C) = \alpha_4^r$ for some $\alpha_1, \ldots, \alpha_4$ that will be chosen in the correctness analysis. Let the corresponding LLRs be o_1, o_2, o_3, o_4 (note that $o_1, o_3 > 0$ and $o_2, o_4 < 0$). We also insert two not-equal gadgets observed by OBS: One of them is put between EVAL and B and the other one between B and C. The overall construction is illustrated in Figure 8.

We are reducing to the problem of computing the action of agent OBS at its significant time $t = 2$. Note that OBS observes all gadgets in the graph, and only gadgets. In particular, OBS directly infers the signals of all auxiliary agents in the gadgets (so these signals are indirectly encoded in the reduction), but the same cannot be said about the private signals $\mathcal{S}(x_1), \ldots, \mathcal{S}(x_N)$ of variable agents. The observation history $H(\mathsf{OBS}, 2)$ is naturally determined by specifications of the gadgets.

Figure 8: 3-SAT reduction for $\phi(x) = \eta_1 \wedge \ldots \wedge \eta_M$.

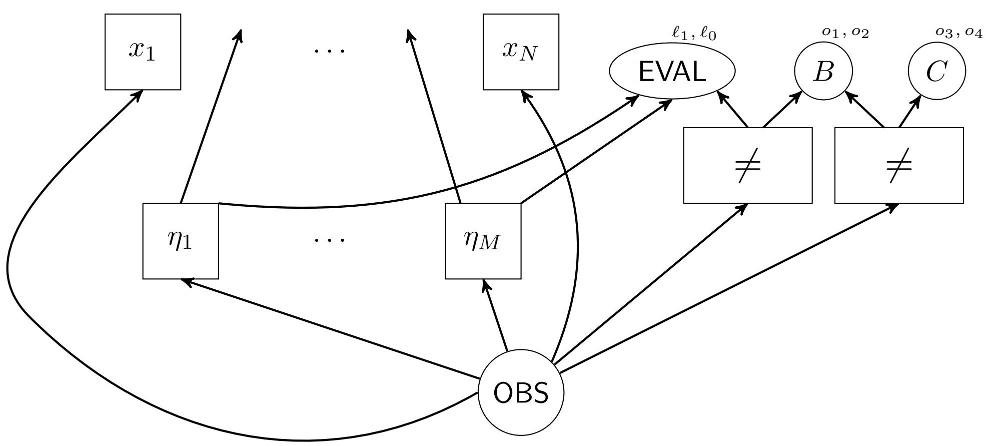

Analysis of the reduction As a preliminary matter, the reduction indeed produces an instance of polynomial size: The size of the graph is $O(N + M)$ and the probabilities of private signals satisfy

$$\exp(-O(N)) \leq p_{\theta_0}(u) \leq 1 - \exp(-O(N)) .$$

We inspect the construction to understand which private signal configurations are consistent with the observation history of agent OBS. Let $\boldsymbol{w} = (w_1, \ldots, w_N)$ be an assignment to the 3-SAT formula. We will say that a signal configuration is consistent with assignment $\boldsymbol{w}$ if it satisfies $\mathcal{S}(x_i) = w_i$ for every variable agent x_i. First, let us summarize what the information available to agent OBS at time $t = 2$.

- Agent OBS observes N variable gadgets. Each of those is a counting gadget with equivalence, itself consisting of two threshold gadgets. The observation history of OBS determines the

private signals of all auxiliary agents in those gadgets. It also ensures that $\mathcal{S}(x_i) \neq \mathcal{S}(\neg x_i)$ for every variable gadget x_i in any consistent signal configuration.
Furthermore, by Claim 10, for each assignment $\boldsymbol{w}$ there is exactly one configuration $\boldsymbol{s}(\boldsymbol{w})$ of private signals of variable gadget agents that is consistent both with $\boldsymbol{w}$ and observation history. Letting $\boldsymbol{S}$ be a random variable consisting of those signals, formula (7) and conditional independence of private signals imply

$$\Pr\left[\boldsymbol{S} = \boldsymbol{s}(\boldsymbol{w}) \wedge \theta = \theta_0\right] = \beta = \beta(N, M) ,$$

where β does not depend[2] on the assignment or on the state of the world θ.
- Agent OBS also observes M clause gadgets, each of them being a counting gadget without equivalence. The gadget corresponding to a clause consisting of literals z_i, z_j, z_k ensures that $\mathcal{S}(z_i) + \mathcal{S}(z_j) + \mathcal{S}(z_k) + \mathcal{S}(\mathsf{EVAL}) \geq 1$.
- Finally, it observes two not-equal gadgets, ensuring $\mathcal{S}(\mathsf{EVAL}) \neq \mathcal{S}(B)$ and $\mathcal{S}(B) \neq \mathcal{S}(C)$.

This leads us to the following:

Claim 12
- *For every assignment $\boldsymbol{w} = (w_1, \ldots, w_N)$, there exists exactly one signal configuration consistent with $\boldsymbol{w}$ and observation history with $\mathcal{S}(\mathsf{EVAL}) = 1, \mathcal{S}(B) = 0, \mathcal{S}(C) = 1$.*
- *For every* satisfying *assignment $\boldsymbol{w}$, there also exists exactly one consistent signal configuration with $\mathcal{S}(\mathsf{EVAL}) = 0, \mathcal{S}(B) = 1, \mathcal{S}(C) = 0$.*
- *There are no other consistent configurations.*

As a next step, we compare the LLRs of configurations corresponding to different assignments. To this end, we let the quantity $P(\boldsymbol{w}, s_0, \theta_0)$ be the a priori probability that private signals are in the consistent configuration corresponding to assignment $\boldsymbol{w}$, $\mathcal{S}(\mathsf{EVAL}) = s_0$ and $\theta = \theta_0$ (note that this is a different definition than given in (3)). Furthermore, we set $P(\boldsymbol{w}, \theta_0) := P(\boldsymbol{w}, 0, \theta_0) + P(\boldsymbol{w}, 1, \theta_0)$.

Keeping in mind the summary of knowledge of agent OBS and using Claims 9 and 10 together with conditional independence of signals we conclude that for any assignment $\boldsymbol{w}$:

$$P(\boldsymbol{w}, 1, \mathsf{T}) = q \cdot 0.9 \cdot \alpha_1^r \cdot (1 - \alpha_3^r) ,$$
$$P(\boldsymbol{w}, 1, \mathsf{F}) = q \cdot 0.4 \cdot (1 - \alpha_2^r) \cdot \alpha_4^r$$

for some $q(N, M) > 0$ that does not depend on $\boldsymbol{w}$. On the other hand, for any satisfying assignment $\boldsymbol{w}$ we additionally have

$$P(\boldsymbol{w}, 0, \mathsf{T}) = q \cdot 0.1 \cdot (1 - \alpha_1^r) \cdot \alpha_3^r ,$$
$$P(\boldsymbol{w}, 0, \mathsf{F}) = q \cdot 0.6 \cdot \alpha_2^r \cdot (1 - \alpha_4^r) .$$

Each of those expressions is a product of four terms. The value q corresponds to the probabilities of signals in variable agents and auxiliary agents in the gadgets. The other terms arise from private signals of, respectively, EVAL, B and C.

2. This is the only place in the reduction where we used the equivalence version of the counting gadget. The point is that the observer agent OBS learns that out of $2N$ "literal agents" exactly N received positive signals. Normally this would imply a shift in belief of OBS, which we avoid by using the equivalence gadget.
In contrast, later on we will be trying to differentiate belief of OBS based on if an assignment is satisfying or not. Intuitively, since we will want to preserve, or even amplify, this difference in belief, there will be no further need for the equivalence gadget.

We choose $\alpha_3 := 0.9$, $\alpha_2 := \alpha_4 := 0.6$, $\alpha_1 := 0.4$ and note that our choice of $r = \gamma N$ for large enough γ ensures that we can estimate[3]

$$P(\boldsymbol{w}, 1, \mathsf{T}) \in q \cdot 0.4^r \cdot \left(1 \pm \frac{1}{200N}\right)^r , \tag{8}$$

$$P(\boldsymbol{w}, 1, \mathsf{F}) \in q \cdot 0.6^r \cdot \left(1 \pm \frac{1}{200N}\right)^r , \tag{9}$$

and, for satisfying assignments,

$$P(\boldsymbol{w}, 0, \mathsf{T}) \in q \cdot 0.9^r \cdot \left(1 \pm \frac{1}{200N}\right)^r . \tag{10}$$

$$P(\boldsymbol{w}, 0, \mathsf{F}) \in q \cdot 0.6^r \cdot \left(1 \pm \frac{1}{200N}\right)^r . \tag{11}$$

This in turn implies that for a satisfying assignment we have

$$P(\boldsymbol{w}, \mathsf{T}) \in q \cdot 0.9^r \cdot \left(1 \pm \frac{1}{100N}\right)^r , P(\boldsymbol{w}, \mathsf{F}) \in q \cdot 0.6^r \cdot \left(1 \pm \frac{1}{100N}\right)^r , \tag{12}$$

and for an unsatisfying one

$$P(\boldsymbol{w}, \mathsf{T}) \in q \cdot 0.4^r \cdot \left(1 \pm \frac{1}{100N}\right)^r , P(\boldsymbol{w}, \mathsf{F}) \in q \cdot 0.6^r \cdot \left(1 \pm \frac{1}{100N}\right)^r . \tag{13}$$

Accordingly, if the formula ϕ has a satisfying assignment $\boldsymbol{w}^*$, it must be that the belief of agent OBS at the significant time $t = 2$ can be bounded by

$$1 - \mu(\mathsf{OBS}) = \frac{\sum_{\boldsymbol{w}\in\{0,1\}^N} P(\boldsymbol{w}, \mathsf{F})}{\sum_{\boldsymbol{w}\in\{0,1\}^N} P(\boldsymbol{w}, \mathsf{F}) + P(\boldsymbol{w}, \mathsf{T})} \le \frac{\sum_{\boldsymbol{w}\in\{0,1\}^N} P(\boldsymbol{w}, \mathsf{F})}{P(\boldsymbol{w}^*, \mathsf{T})} \le \frac{2^N \cdot 0.61^r}{0.89^r} \le 0.69^r . \tag{14}$$

At the same time, this probability can be lower bounded as

$$1 - \mu(\mathsf{OBS}) \ge \frac{P(\boldsymbol{w}^*, \mathsf{F})}{\sum_{\boldsymbol{w}\in\{0,1\}^N} P(\boldsymbol{w}, \mathsf{T}) + P(\boldsymbol{w}, \mathsf{F})} \ge \frac{0.59^r}{2^{N+1} \cdot 0.91^r} \ge 0.64^r . \tag{15}$$

If the formula ϕ is not satisfiable, a simpler computation taking into account only equation (13) gives

$$\mu(\mathsf{OBS}) \in [0.64^r, 0.69^r] . \tag{16}$$

Hence, $\mu(\mathsf{OBS}) = 1 - \exp(-\Theta(N))$ if ϕ is satisfiable and $\mu(\mathsf{OBS}) = \exp(-\Theta(N))$ otherwise. ■

3. In order to facilitate the proof of Theorem 3, the bounds below are slightly better than needed here.

Handling of ties

Remark 13 *There are some results and proofs about opinion exchange models that are sensitive to the tie-breaking rule chosen (see, e.g., Example 3.46 in Mossel and Tamuz (2017)). We claim that the reduction described above (as well as other reductions in this paper) does not suffer from this problem.*

Ideally, we would like to say that ties never arise in signal configurations that are consistent with inputs to the reduction. This is seen to be true by inspection, with the following exception: Agents that do not receive private signals are indifferent about the state of the world until their significant time. We made this choice to simplify the exposition. Since significant times are common knowledge, no agent places any weight on others' actions before their significant time (regardless of the tie-breaking rule used), and the analysis of the reduction is not affected in any way by this fact.

That being said, the ties could be avoided altogether. For example, we could introduce an agent EPS *that is observed by everyone else at time* $t = 0$*, indicating the action* $\mathcal{A}(\mathsf{EPS}) = \mathsf{T}$ *and private signal* $\mathcal{S}(\mathsf{EPS}) = 1$ *corresponding to the LLR* $l_1(\mathsf{EPS}) = \varepsilon$ *for a small constant* $\varepsilon > 0$*. Since LLRs arising in the analysis of our reduction are always bounded away from zero,* ε *can be made small enough so that the agent* EPS *does not affect other agents' actions at their significant times. This almost takes care of the problem, except for the agents without private signals at time* $t = 0$ *(since they will acquire information from* EPS *only at time* $t = 1$*). This can be solved by giving each such agent* u *an informative private signal with LLRs, say,*

$$\ell_1(u) = -\ell_0(u) = \frac{\varepsilon}{100|V|} \; .$$

In that case u *will output an action corresponding to its private signal at time* $t = 0$*, but its belief due to private signal (and signals of all other non-informative agents that* u *observes) will become dominated by the belief of* EPS *at time* $t = 1$*.*

Appendix B. PSPACE-hardness: Proof of Theorem 3

The reduction Recall our formula

$$\Phi = Q_K \boldsymbol{x}_K \cdots \exists \boldsymbol{x}_1 : \phi(\boldsymbol{x}_K, \ldots, \boldsymbol{x}_1) \; .$$

The reduction is defined inductively, with the overall structure illustrated in Figure 10. First, we construct a network identical to the one used used in the 3-SAT reduction for the formula $\phi(\boldsymbol{x}_K, \ldots, \boldsymbol{x}_1)$ (i.e., as if all variables were existential). We call the observer agent OBS_1 and introduce one difference: OBS_1 additionally directly observes all variable agents in variable blocks $\boldsymbol{x}_K, \ldots, \boldsymbol{x}_2$.

Next, for each $1 < i \le k$ we place two agents B_i and C_i with private signals according to probabilities $p_{\mathsf{T}}(B_i) := 1 - \alpha_1^r$, $p_{\mathsf{F}}(B_i) := \alpha_2^r$, $p_{\mathsf{T}}(C_i) := 1 - \alpha_3^r$, $p_{\mathsf{F}}(C_i) := \alpha_4^r$. The parameter r is chosen in the same way as in the 3-SAT reduction, i.e., $r = \gamma \cdot N$ for some absolute γ big enough. The α_j values depend on the parity of i and are provided in Table[4] 1.

4. These values are related to, but different from $\alpha_1, \ldots, \alpha_4$ in Section A. This is because the computations in the base case are different than in the inductive step.

Table 1: Values of α_j for even and odd i.

	even i	odd i
α_1	$\frac{4}{9} \cdot 0.9$	0.9
α_2	0.9	$\frac{4}{9} \cdot 0.9$
α_3	0.9	$\frac{4}{9} \cdot 0.9$
α_4	$\frac{4}{9} \cdot 0.9$	0.9
δ for "large" threshold	$0.2r$	r
δ for "small" threshold	$-r$	$-0.2r$

We place a not-equal gadget between B_i and C_i. We would also like to place a not-equal gadget between OBS_{i-1} and B_i. More precisely, we want a gadget that will convey that relevant actions are different: $\mathcal{A}(\mathsf{OBS}_{i-1}) \neq \mathcal{A}(B_i)$. We cannot use the standard not-equal gadget directly, since OBS_{i-1} receives more complicated information than a single private signal. We now describe how to overcome this difficulty, with an illustration in Figure 9.

We put in place a gadget between OBS_{i-1} and B_i and we want it to be functionally equivalent to a not-equal gadget. We will call it a *modified not-equal gadget*. It consists of two *modified threshold gadgets*. One of those gadgets ensures that $\mathcal{A}(\mathsf{OBS}_{i-1}) \neq \mathsf{T}$ or $\mathcal{A}(B_i) \neq \mathsf{T}$ (we will call it a "large" threshold), and the other one ensures that $\mathcal{A}(\mathsf{OBS}_{i-1}) \neq \mathsf{F}$ or $\mathcal{A}(B_i) \neq \mathsf{F}$ (this is a "small" threshold). Of course the conjunction of those two guarantees is that $\mathcal{A}(\mathsf{OBS}_{i-1}) \neq \mathcal{A}(B_i)$. Since the analysis of two threshold gadgets is symmetric, we describe only the large threshold.

We call the main agent in the modified threshold gadget T_i (it is an equivalent of B in Figure 1). Agent T_i:

- Does not receive a private signal.
- Observes agents OBS_{i-1} and B_i.
- Additionally observes all agents that OBS_{i-1} observes.
- *Except* that it does not observe variable agents in variable block $\boldsymbol{x}_i$.

The significant time of agent OBS_{i-1} is $t = 2i-2$ and we set the significant time of T_i to $t = 2i-1$.

Furthermore, we place two more agents T_i' and T_i'' corresponding to agents C and D in Figure 1. They both receive private signals and the values of those signals are hardcoded in the reduction. Agent T_i' is observed by OBS_i and T_i, and broadcasts LLR $-\delta$. Agent T_i'' is observed only by OBS_i and broadcasts LLR δ. We still need to define the threshold value δ. This is not immediate, since we only have bounds (14)–(16) on beliefs of agent OBS_{i-1}, but it can be done. Precise values for both large and small thresholds are given in Table 1. For the large threshold we specify in the observation history that at the significant time $\mathcal{A}(T_i) = \mathsf{F}$, while for the small threshold we specify the opposite $\mathcal{A}(T_i) = \mathsf{T}$.

Finally, we place an agent OBS_i that observes the same agents as OBS_{i-1}, except for variable agents in variable block $\boldsymbol{x}_i$. Note that OBS_i does not directly observe OBS_{i-1}. Additionally, OBS_i observes the not-equal gadget between B_i and C_i and the modified not-equal gadget (meaning T_i, T_i', and T_i'' in modified threshold gadgets) between OBS_{i-1} and B_i. OBS_i does not receive a private signal, and its significant time is $t = 2i$.

This concludes the definition of the reduction. We show hardness for the computation of agent OBS_K at time $t = 2K$. Again, since this agent observes only gadgets, its observation history is

Figure 9: Modified threshold agent illustrated on case $i = K = 2$ and formula $\Phi = \forall \boldsymbol{y} \exists \boldsymbol{x} : \phi(\boldsymbol{y}, \boldsymbol{x}) = 1$. The gadget consists of agents T_2, T_2' and T_2''. These three agents serve the role of B, C and D from Figure 1 and are all observed by agent OBS_2 (see Figure 10). The gadget implements "not-equal" behavior between agents OBS_1 and B_2.

The red arrow emphasizes that agent OBS_1 directly observes variable agents associated with $\boldsymbol{y}$. Some significant times and LLRs are shown.

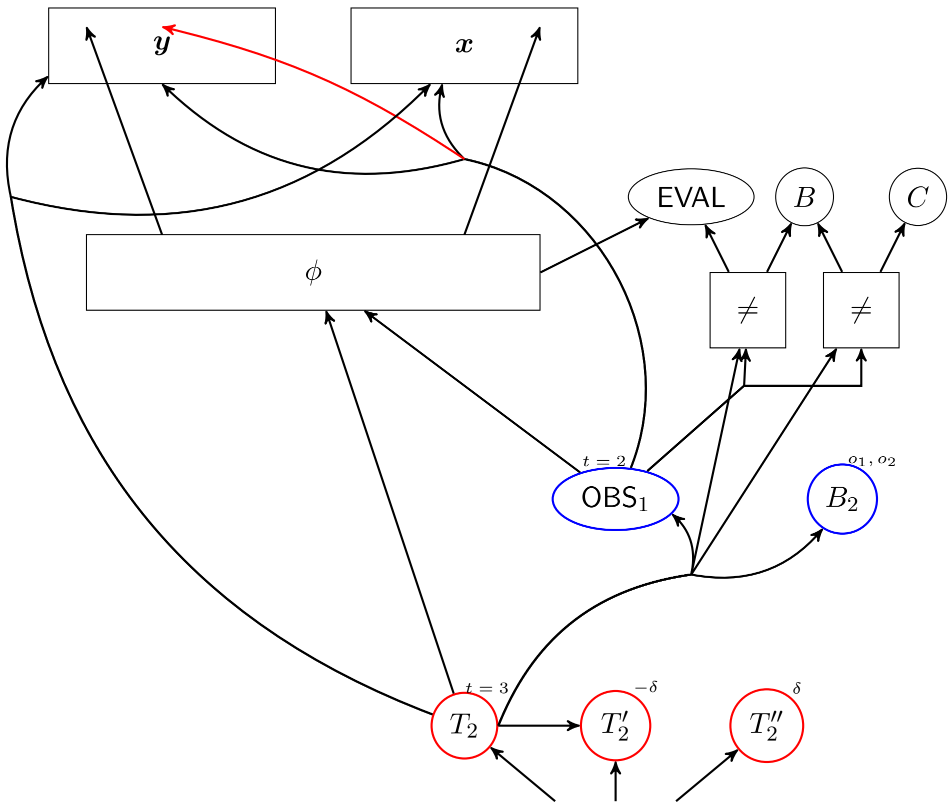

naturally determined by the semantics of the gadgets. We will show that the truth value of formula Φ reduces to distinguishing between $\mu(\mathsf{OBS}_K) \approx 1$ and $\mu(\mathsf{OBS}_K) \approx 0$ and, by implication, $\mathcal{A}(\mathsf{OBS}_K) = \mathsf{T}$ and $\mathcal{A}(\mathsf{OBS}_K) = \mathsf{F}$.

Analysis: Preliminaries To start with, we note that the i-th stage of the inductive definition adds $O(i(N + M))$ new agents (remembering that there are dummy agents that are not shown in the figures). Consequently, the total number of agents is $O(K^2(N + M)) \leq O(N^2(N + M))$. Furthermore, the signal probabilities satisfy (2) by design.

To analyse the belief of agent OBS_K, we need to start with more notation and definitions. For $i > 1$ and a partial assignment to variable blocks $\boldsymbol{w} := (\boldsymbol{w}_K, \ldots, \boldsymbol{w}_i)$, let $\Phi_{\boldsymbol{w}}$ be the formula

$$\Phi_{\boldsymbol{w}} := Q_{i-1}\boldsymbol{x}_{i-1} \cdots \exists \boldsymbol{x}_1 : \phi(\boldsymbol{w}_k, \ldots, \boldsymbol{w}_i, \boldsymbol{x}_{i-1}, \ldots, \boldsymbol{x}_1) \ ,$$

i.e., the original formula with "hard-coded" values of $\boldsymbol{w}$.

Figure 10: Schematic representation of the network in case $K = 2$ for formula $\Phi = \forall \boldsymbol{y} \exists \boldsymbol{x} : \phi(\boldsymbol{y}, \boldsymbol{x}) = 1$. The agents and gadgets added in the inductive definition for $i = 2$ are marked in blue. For clarity, edges from the modified not-equal gadget (cf. Figure 9, here marked with an exclamation point) are not shown.

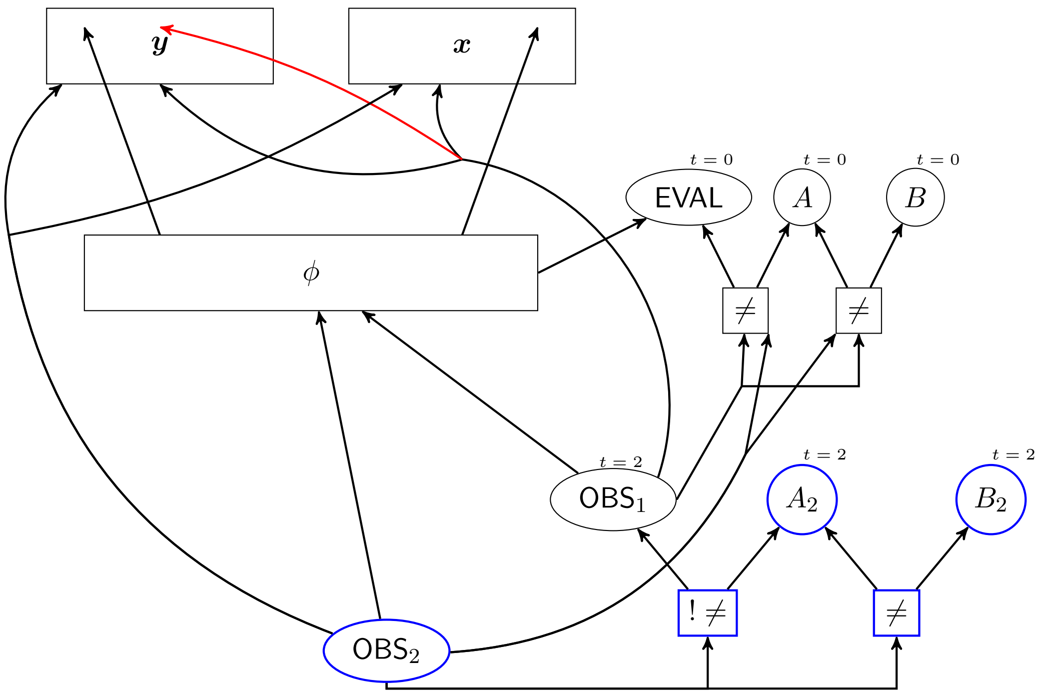

Let G_i be the part of the network consisting of all agents created up to the i-th step of our inductive definition. Therefore, $G = G_K \supseteq \ldots \supseteq G_1$. Note that all variable and clause gadgets are already present in G_1. Compared to G_{i-1}, the subgraph G_i additionally contains agents OBS_i (observing private signals of $\boldsymbol{x}_K, \ldots, \boldsymbol{x}_{i+1}$), B_i and C_i, and two modified not-equal gadgets.

The network G was defined so that all actions of agents in G_i depend only on private signals of agents in G_i. Furthermore, the belief $\mu(\mathsf{OBS}_i)$ depends only on private signals of variable agents $\boldsymbol{x}_K, \ldots, \boldsymbol{x}_{i+1}$ and observations of gadgets by OBS_i (with the latter determined by the reduction, since OBS_K observes all those gadgets as well).

We now need a careful definition in a similar vein to $P(\boldsymbol{w}; \theta_0)$ from the 3-SAT reduction. Given i, $1 \le i \le K$, an assignment $(\boldsymbol{v}, \boldsymbol{w}) := (\boldsymbol{v}_K, \ldots, \boldsymbol{v}_{i+1}, \boldsymbol{w}_i, \ldots, \boldsymbol{w}_1)$, as well as $\theta_0 \in \{\mathsf{T}, \mathsf{F}\}$ we let $P_i(\boldsymbol{v}, \boldsymbol{w}; \theta_0)$ be the probability that all of the following hold:

1. For all gadgets observed by the agent OBS_i, OBS_i observed the actions given by the reduction.
2. The assignment determined by the private signals of the variable agents is equal to $(\boldsymbol{v}, \boldsymbol{w})$.
3. The state of the world is $\theta = \theta_0$.

One checks that $P_i(\boldsymbol{v}, \boldsymbol{w}; \theta_0)$ depends only on private signals in G_i. To gain intuition, the reader is invited to convince oneself that, provided that the modified not-equal gadget ensures $\mathcal{A}(\mathsf{OBS}_{i-1}) \neq \mathcal{A}(B_i)$ (we still need to prove that), $P_i(\boldsymbol{v}, \boldsymbol{w}; \theta_0)$ is always a sum over probabilities of one (if $\phi(\boldsymbol{v}, \boldsymbol{w}) = 0$) or two (in case $\phi(\boldsymbol{v}, \boldsymbol{w}) = 1$) signal configurations on G_i.

Finally, given $\boldsymbol{v}$ and $\alpha, \varepsilon \in (0,1)$ we will say that state of the world θ_0 is α-likely with error ε if both

$$\exists \boldsymbol{w} : P_i(\boldsymbol{v}, \boldsymbol{w}; \theta_0) \geq \alpha^r \cdot (1-\varepsilon)^r \ ,$$
$$\forall \boldsymbol{w} : P_i(\boldsymbol{v}, \boldsymbol{w}; \theta_0) \leq \alpha^r \cdot (1+\varepsilon)^r \ .$$

The analysis proceeds by induction on the block number i, with a two-part invariant we need to maintain. The first part says that, letting $\varepsilon := \frac{i}{100N}$, there exists some $\beta := \beta(i) \in (0,1)$ such that for every partial assignment $\boldsymbol{v} := (\boldsymbol{v}_K, \dots, \boldsymbol{v}_{i+1})$:

1. If i is odd and $\Phi_{\boldsymbol{v}}$ is true, then T is β-likely with error ε and F is $\frac{2}{3}\beta$-likely with error ε.
2. If i is odd and $\Phi_{\boldsymbol{v}}$ is false, then T is $\frac{4}{9}\beta$-likely with error ε and F is $\frac{2}{3}\beta$-likely with error ε.
3. Symmetrically, if i is even and $\Phi_{\boldsymbol{v}}$ is true, then T is $\frac{2}{3}\beta$-likely with error ε and F is $\frac{4}{9}\beta$-likely with error ε.
4. If i is even and $\Phi_{\boldsymbol{v}}$ is false, then T is $\frac{2}{3}\beta$-likely with error ε and F is β-likely with error ε.

The second part of the invariant states that whenever $\Phi_{\boldsymbol{v}}$ is true, the belief of agent OBS_i satisfies $1-\mu(\mathsf{OBS}_i) \in [0.64^r, 0.69^r]$. Similarly, if $\Phi_{\boldsymbol{v}}$ is false, then this belief satisfies $\mu(\mathsf{OBS}_i) \in [0.64^r, 0.69^r]$. Note that this part applied to $i = K$ implies the last bullet point in the statement of Theorem 3, with $\mu(\mathsf{OBS}_K)$ being within $\exp(-\Theta(N))$ distance to either zero or one.

Base case To establish the base case $i = 1$ one has to go through the proof in Section A and convince themselves that the analysis stays valid even when the agent OBS directly observes variable agents $\boldsymbol{v}_K, \dots, \boldsymbol{v}_2$. Then, the first invariant is established with

$$\beta(1) := q^{1/r} \cdot 0.9 \ ,$$

where q is the value featured in equations (12)-(13). For example, $\Phi_{\boldsymbol{v}}$ being true means that the respective 3-CNF formula $\phi_{\boldsymbol{v}}(\boldsymbol{x}_1)$ is satisfiable. Taking a satisfying assignment $\boldsymbol{w}$, we get by (12)

$$P_1(\boldsymbol{v}, \boldsymbol{w}; \mathsf{T}) = P(\boldsymbol{w}; \mathsf{T}) \geq q \cdot 0.9^r \cdot (1-\varepsilon)^r = \beta^r \cdot (1-\varepsilon)^r \ ,$$
$$P_1(\boldsymbol{y}, \boldsymbol{w}; \mathsf{F}) = P(\boldsymbol{w}; \mathsf{F}) \geq q \cdot 0.6^r \cdot (1-\varepsilon)^r = \left(\frac{2}{3}\beta\right)^r \cdot (1-\varepsilon)^r \ .$$

On the other hand, by (12) and (13), for every assignment $\boldsymbol{w}$, satisfying or not, we have

$$P_1(\boldsymbol{v}, \boldsymbol{w}; \mathsf{T}) \leq \max\left(q \cdot 0.9^r \cdot (1+\varepsilon)^r, q \cdot 0.4^r \cdot (1+\varepsilon)^r\right) \leq \beta^r \cdot (1+\varepsilon)^r \ ,$$
$$P_1(\boldsymbol{v}, \boldsymbol{w}; \mathsf{F}) \leq q \cdot 0.6^r \cdot (1+\varepsilon)^r = \left(\frac{2}{3}\beta\right)^r \cdot (1+\varepsilon)^r \ .$$

A similar computation gives the first invariant in case $\Phi_{\boldsymbol{v}}$ is false, this time using only (13). The second invariant is a direct consequence of equations (14)-(16).

Induction step We will analyze only even i, since the other case is analogous. Fix some $\boldsymbol{v} = (\boldsymbol{v}_K, \dots, \boldsymbol{v}_{i+1})$. In the following we assume that all actions observed in gadgets are as given by the reduction and that private signals for the initial blocks of variables are given by $\boldsymbol{v}$. Let us call private signal configurations on G_{i-1} that satisfy those conditions *consistent*.

In this setting, every assignment of private signals $\boldsymbol{v}_i$ in variable block $\boldsymbol{x}_i$ determines the action $\mathcal{A}(\mathsf{OBS}_{i-1})$ and, by the second invariant, $\mathcal{A}(\mathsf{OBS}_{i-1}) = \mathsf{T}$ if and only if the formula $\Phi_{\boldsymbol{v},\boldsymbol{v}_i}$ is true. Accordingly, we divide consistent configurations into "T-configurations" and "F-configurations".

Our first objective is to show that the modified not-equal gadget (cf. Figure 10) ensures that $\mathcal{A}(\mathsf{OBS}_{i-1}) \neq \mathcal{A}(B_i)$. Let T_i be the main agent in a modified threshold gadget between OBS_{i-1} and B_i (cf. Figure 9). At its significant time $t = 2i - 1$, agent T_i observed everything that agent OBS_{i-1} observed except for the assignment $\boldsymbol{v}_i$. It also observed the action $\mathcal{A}(\mathsf{OBS}_{i-1}) = \theta_0$. Therefore, the signal configurations on G_{i-1} consistent with observations of T_i are exactly the θ_0-configurations. We let

$$p_{\mathsf{OBS}}(\theta_0) := \mathrm{E}\left[\mu(\mathsf{OBS}_{i-1})\right] ,$$

where the expectation is over all θ_0-configurations. By the second invariant, $p_{\mathsf{OBS}}(\mathsf{F})$ is an average over terms $\mu(\mathsf{OBS}_{i-1})$ such that each term satisfies $\mu(\mathsf{OBS}_{i-1}) \in [0.64^r, 0.69^r]$. Therefore, we have

$$0.64^r \le p_{\mathsf{OBS}}(\mathsf{F}) \le 0.69^r .$$

Similarly,

$$0.64^r \le 1 - p_{\mathsf{OBS}}(\mathsf{T}) \le 0.69^r .$$

Let $m(\theta_0) := \ln \frac{p_{\mathsf{OBS}}(\theta_0)}{1-p_{\mathsf{OBS}}(\theta_0)}$. We check that

$$m(\mathsf{T}) \in [0.37r, 0.45r],\ m(\mathsf{F}) \in [-0.45r, -0.37r] . \tag{17}$$

We interpret $m(\theta_0)$ as the LLR of agent T_i based only on observations from G_{i-1}, excluding its observations of B_i and T'_i. On the other hand, T_i directly observes B_i and T'_i, which receive private signals. The LLRs of B_i are given by

$$\ell_1(B_i) = \ln \frac{1-\alpha_1^r}{\alpha_2^r} = \ln \frac{1-\left(\frac{4}{9}\cdot 0.9\right)^r}{0.9^r} \in [0.1r, 0.11r] \tag{18}$$

$$\ell_0(B_i) = \ln \frac{\alpha_1^r}{1-\alpha_2^r} = \ln \frac{\left(\frac{4}{9}\cdot 0.9\right)^r}{1-0.9^r} \in [-0.92r, -0.91r] . \tag{19}$$

The LLR of T'_i is equal to $-\delta$ given by Table 1 depending on which modified threshold gadget we are considering. If it is the "large" threshold, ensuring that $\mathcal{A}(\mathsf{OBS}_{i-1}) = \mathsf{F}$ or $\mathcal{A}(B_i) = \mathsf{F}$, then $\delta = 0.2r$. If it is the "small" threshold, ensuring $\mathcal{A}(\mathsf{OBS}_{i-1}) = \mathsf{T}$ or $\mathcal{A}(B_i) = \mathsf{T}$, then $\delta = -r$.

To sum up, we can see agent T_i as receiving information from three sources: subnetwork G_{i-1}, agent B_i and agent T'_i. Furthermore, signals from those three sources are independent. Therefore, we can write the LLR of agent T_i at its significant time as

$$L = m(\theta_1) + \ell_b(B_i) - \delta , \tag{20}$$

where $\theta_1 := \mathcal{A}(\mathsf{OBS}_{i-1})$ and $b := \mathcal{S}(B_i)$.

Let us now focus on the large threshold. Recall that in this case we specified the action at the significant time as $\mathcal{A}(T_i) = \mathsf{F}$, which by (20) means

$$m(\theta_1) + \ell_b(B_i) < 0.2r . \tag{21}$$

But the bounds (17)-(19) imply that (21) holds if and only if $\theta_1 = \mathsf{F}$ or $b = 0$, i.e., if and only if $\mathcal{A}(\mathsf{OBS}_{i-1}) = \mathsf{F}$ or $\mathcal{A}(B_i) = \mathsf{F}$.

Similarly, consider agent T_i in the small threshold gadget. Since we defined its action at the significant time as $\mathcal{A}(T_i) = \mathsf{T}$, similarly to (21) we have

$$m(\theta_1) + \ell_b(B_i) > -r \,. \tag{22}$$

Again checking (22) against (17)-(19), we see that (22) holds if and only if $\theta_1 = \mathsf{T}$ or $b = 1$, i.e., if and only if $\mathcal{A}(\mathsf{OBS}_{i-1}) = \mathsf{T}$ or $\mathcal{A}(B_i) = \mathsf{T}$.

Therefore, by observing two agents T_i in the modified threshold agents, agent OBS_i learns that $\mathcal{A}(\mathsf{OBS}_{i-1}) \neq \mathcal{A}(B_i)$. Furthermore, recall that OBS_i observes everything that OBS_{i-1} observes except for signals in variable block $\boldsymbol{x}_i$. As a result, we have the following claim:

Claim 14 *Any configuration in G_{i-1} consistent with observations of OBS_{i-1} uniquely extends to a configuration in G_i consistent with observations of OBS_i such that $\mathcal{A}(\mathsf{OBS}_{i-1}) \neq \mathcal{A}(B_i)$ and $\mathcal{A}(B_i) \neq \mathcal{A}(C_i)$.*

What agent OBS_i does not know compared to OBS_{i-1} is the assignment $\boldsymbol{v}_i$ to variable block $\boldsymbol{x}_i$. Therefore, the belief of OBS_i is an average over all possible values of $\boldsymbol{v}_i$. What is more, a term in this average for fixed $\boldsymbol{v}_i$ is a product of the following independent probabilities (beliefs):

- The belief coming from G_{i-1}, which is equal to the belief of OBS_i, for assignment $(\boldsymbol{v}, \boldsymbol{v}_i)$.
- The beliefs coming from auxiliary agents in the modified not-equal gadget between OBS_{i-1} and B_i and the not-equal gadget between B_i and C_i. By design, these agents do not affect the overall belief of OBS_i.
- The beliefs of B_i and C_i.

By inductive assumption, if $\Phi_{\boldsymbol{v},\boldsymbol{v}_i} \in \mathrm{TQBF}$, then $\mathcal{A}(\mathsf{OBS}_{i-1}) = \mathsf{T}$, $S(B_i) = 0$ and $S(C_i) = 1$. As a result, for any assignment to all variables $(\boldsymbol{v}, \boldsymbol{v}_i, \boldsymbol{w})$ we can write (cf. Table 1)

$$\begin{aligned} P_i(\boldsymbol{v}, \boldsymbol{v}_i, \boldsymbol{w}; \mathsf{T}) &= P_{i-1}(\boldsymbol{v}, \boldsymbol{v}_i, \boldsymbol{w}; \mathsf{T}) \cdot q \cdot \alpha_1^r \cdot (1 - \alpha_3^r) \\ &\in P_{i-1}(\boldsymbol{v}, \boldsymbol{v}_i, \boldsymbol{w}; \mathsf{T}) \cdot q \cdot \left(0.9 \cdot \frac{4}{9}\right)^r \cdot \left(1 \pm \frac{1}{200N}\right)^r \\ P_i(\boldsymbol{v}, \boldsymbol{v}_i, \boldsymbol{w}; \mathsf{F}) &= P_{i-1}(\boldsymbol{v}, \boldsymbol{v}_i, \boldsymbol{w}; \mathsf{F}) \cdot q \cdot (1 - \alpha_2)^r \cdot \alpha_4^r \\ &\in P_{i-1}(\boldsymbol{v}, \boldsymbol{v}_i, \boldsymbol{w}; \mathsf{F}) \cdot q \cdot \left(0.9 \cdot \frac{4}{9}\right)^r \cdot \left(1 \pm \frac{1}{200N}\right)^r \end{aligned} \tag{23}$$

where $P_{i-1}(\cdot)$, q and α_j factors come from the three sources described in the items above.

On the other hand, if $\Phi_{\boldsymbol{v},\boldsymbol{v}_i} \notin \mathrm{TQBF}$, then $\mathcal{A}(\mathsf{OBS}_{i-1}) = \mathsf{F}$, $S(B_i) = 1$, $S(C_i) = 0$ and, similarly

$$P_i(\boldsymbol{v}, \boldsymbol{v}_i, \boldsymbol{w}; \theta_0) \in P_{i-1}(\boldsymbol{v}, \boldsymbol{v}_i, \boldsymbol{w}; \theta_0) \cdot q \cdot 0.9^r \cdot \left(1 \pm \frac{1}{200N}\right)^r \,. \tag{24}$$

We conclude by establishing both invariants. Since i is even, by induction we know that for some $\beta' = \beta(i-1) \in (0,1)$, if $\Phi_{\boldsymbol{v},\boldsymbol{v}_i}$ is true, then T is β'-likely and F is $\frac{2}{3}\beta'$-likely, and if $\Phi_{\boldsymbol{v},\boldsymbol{v}_i}$ is false, then T is $\frac{4}{9}\beta'$-likely and F is $\frac{2}{3}\beta'$-likely, all with error $\frac{i-1}{100N}$.

Take $\beta := \beta(i) := \beta' \cdot q^{1/r} \cdot 0.9 \cdot \frac{2}{3}$. We want to establish that if $\Phi_{\boldsymbol{v}}$ is true, then T is $\frac{2}{3}\beta$-likely and F is $\frac{4}{9}\beta$-likely, and if $\Phi_{\boldsymbol{v}}$ is false, then T is $\frac{2}{3}\beta$-likely and F is β-likely, all with error $\varepsilon := \frac{i}{100N}$.

Since i is even, the quantifier Q_i is universal and $\Phi_{\boldsymbol{v}}$ is true if and only if $\Phi_{\boldsymbol{v},\boldsymbol{v}_i}$ is true for all $\boldsymbol{v}_i$. To establish the first invariant in this case, we want to show two things. First, for every $\boldsymbol{v}_i, \boldsymbol{w}$ it should hold $P(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w};\mathsf{T}) \le \left(\frac{2}{3}\beta(1+\varepsilon)\right)^r$ and $P(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w};\mathsf{F}) \le \left(\frac{4}{9}\beta(1+\varepsilon)\right)^r$. Second, there should exist $\boldsymbol{v}_i, \boldsymbol{w}$ such that $P(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w};\mathsf{T}) \ge \left(\frac{2}{3}\beta(1-\varepsilon)\right)^r$ and $P(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w};\mathsf{F}) \ge \left(\frac{4}{9}\beta(1-\varepsilon)\right)^r$.

Indeed, by (23) and induction, if $\Phi_{\boldsymbol{v},\boldsymbol{v}_i}$ is true, then for every $\boldsymbol{w}$:

$$
\begin{aligned}
P_i(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w};\mathsf{T}) &\le P_{i-1}(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w};\mathsf{T}) \cdot q \cdot \left(0.9 \cdot \frac{4}{9}\right)^r \left(1 + \frac{1}{200N}\right)^r \\
&\le (\beta')^r \left(1 + \frac{i-1}{100N}\right)^r \cdot q \cdot \left(0.9 \cdot \frac{4}{9}\right)^r \left(1 + \frac{1}{200N}\right)^r \\
&\le \left(\frac{2}{3} \cdot \beta\right)^r (1+\varepsilon)^r , && (25) \\
P_i(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w};\mathsf{F}) &\le \left(\frac{2}{3} \cdot \beta'\right)^r \left(1 + \frac{i-1}{100N}\right)^r \cdot q \cdot \left(0.9 \cdot \frac{4}{9}\right)^r \left(1 + \frac{1}{200N}\right)^r \\
&\le \left(\frac{4}{9} \cdot \beta\right)^r (1+\varepsilon)^r . && (26)
\end{aligned}
$$

At the same time, we can take arbitrary $\boldsymbol{v}_i$ and $\boldsymbol{w}^{\mathsf{T}}, \boldsymbol{w}^{\mathsf{F}}$ that achieve, respectively $P(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w}^{\mathsf{T}};\mathsf{T}) \ge \left(\beta'(1-\frac{i-1}{100N})\right)^r$ and $P(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w}^{\mathsf{F}};\mathsf{F}) \ge \left(\frac{2}{3}\beta'(1-\frac{i-1}{100N})\right)^r$ and see that we have

$$
\begin{aligned}
P_i(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w}^{\mathsf{T}};\mathsf{T}) &\ge \left(\frac{2}{3} \cdot \beta\right)^r (1-\varepsilon)^r , && (27) \\
P_i(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w}^{\mathsf{F}};\mathsf{F}) &\ge \left(\frac{4}{9} \cdot \beta\right)^r (1-\varepsilon)^r , && (28)
\end{aligned}
$$

concluding that T is $\frac{2}{3}\beta$-likely and F is $\frac{4}{9}\beta$ likely with error ε, just as we wanted.

On the other hand, if $\Phi_{\boldsymbol{v}}$ is false, then we have two cases to consider. First, $\Phi_{\boldsymbol{v},\boldsymbol{v}_i}$ can be true, in which case upper bounds (25)-(26) hold. However, we also know that there exists $\boldsymbol{v}_i$ such that $\Phi_{\boldsymbol{v},\boldsymbol{v}_i}$ is false, in which case, using (24), for all $\boldsymbol{w}$:

$$
\begin{aligned}
P_i(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w};\mathsf{T}) &\le P_{i-1}(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w};\mathsf{T}) \cdot q \cdot 0.9^r \cdot \left(1 + \frac{1}{200N}\right)^r \\
&\le \left(\frac{4}{9} \cdot \beta'\right)^r \left(1 + \frac{i-1}{100N}\right)^r \cdot q \cdot 0.9^r \cdot \left(1 + \frac{1}{200N}\right)^r \\
&\le \left(\frac{2}{3} \cdot \beta\right)^r (1+\varepsilon)^r , \\
P_i(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w};\mathsf{F}) &\le \left(\frac{2}{3} \cdot \beta'\right)^r \left(1 + \frac{i-1}{100N}\right)^r \cdot q \cdot 0.9^r \cdot \left(1 + \frac{1}{200N}\right)^r \\
&\le \beta^r \cdot (1+\varepsilon)^r .
\end{aligned}
$$

and similarly there exist $\boldsymbol{w}^{\mathsf{T}}, \boldsymbol{w}^{\mathsf{F}}$ such that

$$
\begin{aligned}
P_i(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w}^{\mathsf{T}};\mathsf{T}) &\ge \left(\frac{2}{3} \cdot \beta\right)^r (1-\varepsilon)^r , \\
P_i(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w}^{\mathsf{F}};\mathsf{F}) &\ge \beta^r \cdot (1-\varepsilon)^r ,
\end{aligned}
$$

establishing that T is $\frac{2}{3}\beta$-likely and F is β-likely with error ε. This concludes the proof of the first invariant

Finally, we need to use a computation similar as in (14) and (15) to check the second invariant. If $\Phi_{\boldsymbol{v}}$ is true, then, since T is $\frac{2}{3}\beta$-likely and F is $\frac{4}{9}\beta$-likely with error $\frac{i}{100N}$,

$$\begin{aligned}
1-\mu(\mathsf{OBS}_i) &= \frac{\sum_{\boldsymbol{v}_i,\boldsymbol{w}} P_i(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w};\mathsf{F})}{\sum_{\boldsymbol{v}_i,\boldsymbol{w}} P_i(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w};\mathsf{T})+P_i(\boldsymbol{v},\boldsymbol{v}_i,\boldsymbol{w};\mathsf{F})}\\
&\le \frac{2^N\beta^r(4/9)^r(1+i/100N)^r}{\beta^r(2/3)^r(1-i/100N)^r} \le \frac{2^N(2/3)^r 1.01^r}{0.99^r} \le (2/3)^r\cdot 1.03^r\cdot 2^N \le 0.69^r\ ,\\
1-\mu(\mathsf{OBS}_i) &\ge \frac{\beta^r(4/9)^r(1-i/100N)^r}{2^{N+1}\beta^r(2/3)^r(1+i/100N)^r} \ge \frac{(2/3)^r 0.99^r}{2^{N+1}1.01^r} \ge (2/3)^r\cdot 0.97^r\cdot 2^N \ge 0.64^r\ .
\end{aligned}$$

A symmetric computation confirms that the second invariant is preserved also when $\Phi_{\boldsymbol{v}}$ is false. ■

Appendix C. Bounded signals: Proof of Theorem 7

One could object that our reduction uses private signal distributions with probabilities that are exponentially close to zero and one. Given that it is a worst-case reduction, with relevant configurations arising with exponentially small probability, we do not think this is a significant issue. In any case, in this section we explain how to modify the proof of Theorem 3 so that it uses only a fixed collection of (say, at most fifty) private signal distributions.

Note that the only agents we need to replace are B and C from the 3-SAT reduction, and B_i, C_i from the induction step in the PSPACE reduction, as well as their associated not-equal gadgets. We sketch the modifications on one example, since other cases are analogous. To this end, take even i and consider B_i, C_i and their not-equal gadgets (cf. Figures 9 and 10).

Going back to the proof of Theorem 3, in particular equations (23)-(24), what we would like to have is that for every consistent configuration on G_{i-1}, there should be a unique way of extending it to a consistent configuration on G_i such that for an assignment $(\boldsymbol{v},\boldsymbol{w})$ and $\theta_0 \in \{\mathsf{F},\mathsf{T}\}$,

$$P_i(\boldsymbol{v},\boldsymbol{w};\theta_0) = P_{i-1}(\boldsymbol{v},\boldsymbol{w};\theta_0)\cdot q\cdot \begin{cases} \alpha_1^r & \text{if } \mathcal{A}(\mathsf{OBS}_{i-1})=\mathsf{T} \text{ and } \theta_0=\mathsf{T},\\ \alpha_4^r & \text{if } \mathcal{A}(\mathsf{OBS}_{i-1})=\mathsf{T} \text{ and } \theta_0=\mathsf{F},\\ \alpha_3^r & \text{if } \mathcal{A}(\mathsf{OBS}_{i-1})=\mathsf{F} \text{ and } \theta_0=\mathsf{T},\\ \alpha_2^r & \text{if } \mathcal{A}(\mathsf{OBS}_{i-1})=\mathsf{F} \text{ and } \theta_0=\mathsf{F},\end{cases}$$

for some $q \in (0,1)$ independent of $(\boldsymbol{v},\boldsymbol{w},\theta_0)$. We are going to achieve this using two independent gadgets corresponding to B_i and C_i. Again, we only sketch the construction for B_i. What we need, then, is to create a gadget that extends every consistent configuration on G_{i-1} to a unique consistent configuration on G_i such that

$$P_i(\boldsymbol{v},\boldsymbol{w};\theta_0) = P_{i-1}(\boldsymbol{v},\boldsymbol{w};\theta_0)\cdot q\cdot \begin{cases} \alpha_1^r & \text{if } \mathcal{A}(\mathsf{OBS}_{i-1})=\mathsf{T}=\theta_0,\\ \alpha_2^r & \text{if } \mathcal{A}(\mathsf{OBS}_{i-1})=\mathsf{F}=\theta_0,\\ 1 & \text{otherwise.}\end{cases} \tag{29}$$

This is achieved as shown in Figure 11. We create an agent H with fixed, arbitrary distribution, say $p_{\mathsf{F}}(H) = 1/4$ and $p_{\mathsf{T}}(H) = 3/4$. Then, we add agents D_j, E_j, F_j for $j = 1, \dots, r$ with private signal distributions

$$p_{\theta_0}(E_j) := p_{\theta_0} \; ,$$
$$p_{\theta_0}(D_j) := p_{\theta_0}(F_j) := q_{\theta_0} \; ,$$

for some (dependent on α_i) values $p_{\mathsf{F}}, p_{\mathsf{T}}, q_{\mathsf{F}}, q_{\mathsf{T}}$ that we will specify shortly.

For each triple D_j, E_j, F_j we also place three not-equal gadgets observed by OBS_i: Respectively, between H and D_j, D_j and E_j, and E_j and F_j. We also create an agent H' with the same signal distribution as H, and a counting gadget with equivalence observed by OBS_i, making sure that $\mathcal{S}(H) + \mathcal{S}(H') = 1$ (this is to get rid of a small distortion in (29) due to the signal of agent H; we will not worry about it from now on). Finally, we place a gadget between OBS_{i-1} and H generalizing the modified not-equal gadget from Theorem 3. This gadget will be observed by OBS_i and we will fill in its details later.

Figure 11: Bounded signals gadget. One out of r parts is shown. The details of the modified not-equal gadget between OBS_{i-1} and H are not shown, and the counting gadget between H and H' is not included.

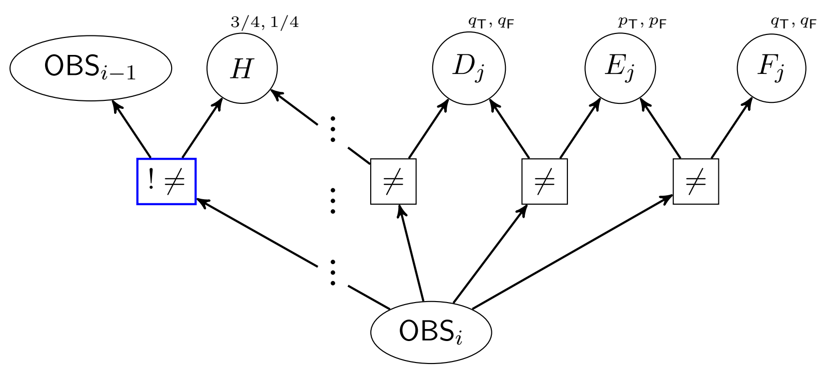

Let us assume for now that the modified not-equal gadget ensures that $\mathcal{A}(\mathsf{OBS}_{i-1}) \neq \mathcal{A}(H)$ in every consistent configuration. Then, since not-equal gadgets guarantee $\mathcal{S}(H) = \mathcal{S}(E_j) \neq \mathcal{S}(D_j) = \mathcal{S}(F_j)$ for every j, we claim that it is not difficult to see that every consistent configuration on G_{i-1} can be uniquely extended to a consistent configuration on G_i such that

$$P_i(\boldsymbol{v}, \boldsymbol{w}; \theta_0) = P_{i-1}(\boldsymbol{v}, \boldsymbol{w}; \theta_0) \cdot q \cdot \begin{cases} \left(q_{\mathsf{T}}^2(1-p_{\mathsf{T}})\right)^r & \text{if } \mathcal{A}(\mathsf{OBS}_{i-1}) = \mathsf{T} \text{ and } \theta_0 = \mathsf{T}, \\ \left(q_{\mathsf{F}}^2(1-p_{\mathsf{F}})\right)^r & \text{if } \mathcal{A}(\mathsf{OBS}_{i-1}) = \mathsf{T} \text{ and } \theta_0 = \mathsf{F}, \\ \left((1-q_{\mathsf{T}})^2 p_{\mathsf{T}}\right)^r & \text{if } \mathcal{A}(\mathsf{OBS}_{i-1}) = \mathsf{F} \text{ and } \theta_0 = \mathsf{T}, \\ \left((1-q_{\mathsf{F}})^2 p_{\mathsf{F}}\right)^r & \text{if } \mathcal{A}(\mathsf{OBS}_{i-1}) = \mathsf{F} \text{ and } \theta_0 = \mathsf{F}. \end{cases}$$

Comparing with (29), we need to find $p_{\mathsf{F}}, p_{\mathsf{T}}, q_{\mathsf{F}}, q_{\mathsf{T}}$ satisfying

$$\frac{q_{\mathsf{T}}^2(1-p_{\mathsf{T}})}{\alpha_1} = q_{\mathsf{F}}^2(1-p_{\mathsf{F}}) = (1-q_{\mathsf{T}})^2 p_{\mathsf{T}} = \frac{(1-q_{\mathsf{F}})^2 p_{\mathsf{F}}}{\alpha_2} \; . \tag{30}$$

Separately comparing and transforming the terms in (30): second with fourth, and then first with third, we get

$$p_{\mathsf{F}} = \frac{\alpha_2 q_{\mathsf{F}}^2}{\alpha_2 q_{\mathsf{F}}^2 + (1-q_{\mathsf{F}})^2}\,, \qquad p_{\mathsf{T}} = \frac{q_{\mathsf{T}}^2}{q_{\mathsf{T}}^2 + \alpha_1(1-q_{\mathsf{T}})^2}\,,$$

which can be substituted into comparison of the first and second term, yielding

$$\frac{q_{\mathsf{T}}^2(1-q_{\mathsf{T}})^2}{q_{\mathsf{T}}^2 + \alpha_1(1-q_{\mathsf{T}})^2} = \frac{q_{\mathsf{F}}^2(1-q_{\mathsf{F}})^2}{\alpha_2 q_{\mathsf{F}}^2 + (1-q_{\mathsf{F}})^2}\,.$$

Taking $q_{\mathsf{T}} := 1-\varepsilon$ for small enough $\varepsilon > 0$, this can be checked to have a solution with $q_{\mathsf{F}} = \varepsilon + O(\varepsilon^2)$, $p_{\mathsf{F}} = \alpha_2\varepsilon^2 + O(\varepsilon^3)$ and $p_{\mathsf{T}} = 1 - \alpha_1\varepsilon^2 + O(\varepsilon^3)$.

We still need to explain how to construct the modified not-equal gadget ensuring that $\mathcal{A}(\mathsf{OBS}_{i-1}) \neq \mathcal{A}(H)$. This is a generalization of the construction in Figure 9 and is shown in Figure 12.

Figure 12: Implementation of the modified not-equal gadget (marked in blue in Figure 11).

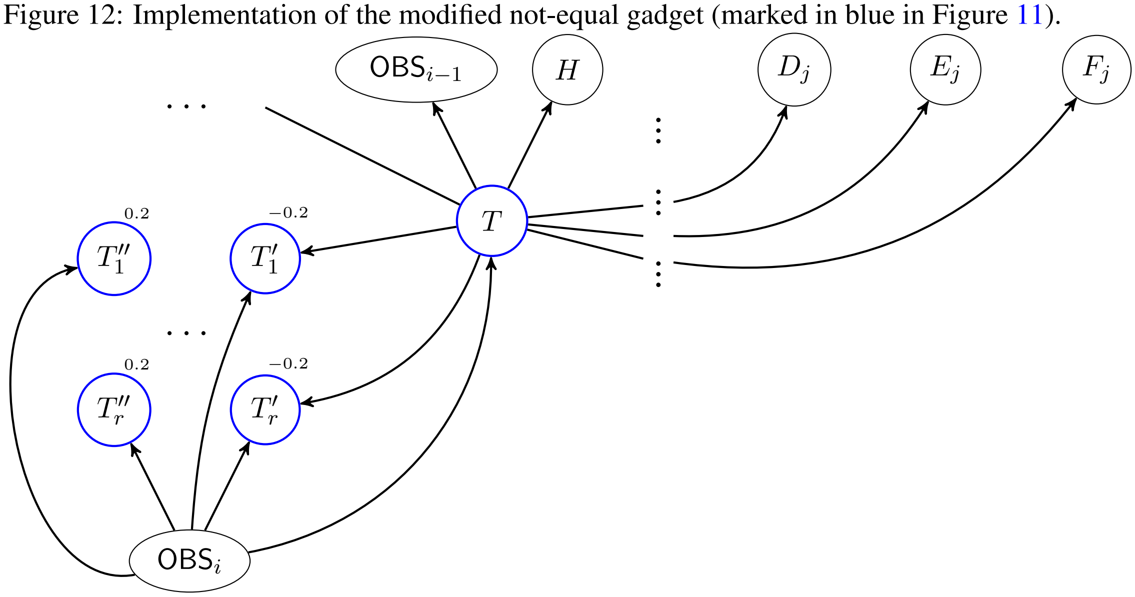

Yet again, it is achieved by combining two threshold gadgets and we focus on one of them. Recall from Table 1 that this threshold was set at $\delta = 0.2r$.

The threshold gadget will have a "counting agent" T and auxiliary agents $T_1', \ldots, T_r'$ and T_1'', $\ldots, T_r''$. Auxiliary agents receive private signals with LLRs $\ell_0(T_j') := -0.2$ and $\ell_1(T_j'') := 0.2$. Agent T observes OBS_{i-1}, as well as other gadgets and agents in the network G_{i-1}, in the same way as agent T_1 in Figure 9. Additionally, it directly observes all agents with private signals in the counting gadget between H and H', as well as all of D_j, E_j and F_j. Finally, it observes $T_1', \ldots, T_r'$. Agent OBS_i observes $T_1', \ldots, T_r', T_1'', \ldots, T_r''$ and T. As expected, we specify that at the significant time OBS_i observes actions $\mathcal{A}(T_j') = \mathsf{F}$, $\mathcal{A}(T_j'') = \mathsf{T}$ and $\mathcal{A}(T) = \mathsf{F}$. Note that we do not need to change the significant time of OBS_i.

Assuming that $\mathcal{A}(\mathsf{OBS}_{i-1}) = \theta_1$ and $\mathcal{A}(F) = \theta_2$, we use the same reasoning as in (20) to compute the LLR $m(\theta_1)$ that agent T can infer from looking at G_{i-1}, another LLR $\ell(\theta_2)$ that can be inferred from looking at H, D_j, E_j, F_j and the LLR $-\delta = -0.2b$ arising from looking at $T'_1, \dots, T'_b$. The bounds on $m(\theta_1)$ are the same as in (17), and as for $\ell(\theta_2)$, from (30) we get, as expected $\ell(\mathsf{F}) = -r \ln \frac{1}{\alpha_1}$ and $\ell(\mathsf{T}) = r \ln \frac{1}{\alpha_2}$.

Since the private signals in these three parts of the graph are conditionally independent, these LLRs can be added up to ensure that $\mathcal{A}(T) = \mathsf{F}$ if and only if

$$m(\theta_1) + \ell(\theta_2) < \delta \ ,$$

which implies, the same as in the proof of Theorem 3, that in a consistent configuration either $\theta_1 = \mathsf{F}$ or $\theta_2 = \mathsf{F}$.

As mentioned, other cases proceed in a similar manner. One difference is that for agents B and C in the base case (3-SAT reduction), EVAL is a simple agent with bounded signal (as opposed to OBS_{i-1}). However, this is only good news for us: We do not need to implement the modified not-equal gadget, since a simple not-equal gadget between C and EVAL suffices. ■

Appendix D. $\#\mathsf{P}$-hardness of revealed beliefs

Revealed belief model: Our result In a natural variant of our model the agents act in exactly the same manner, except that they reveal their full beliefs $\mathcal{A}(u,t) = \mu(u,t)$ rather than just estimates of the state θ. Accordingly, we call it the *revealed belief* model. We suspect that binary action and revealed belief models have similar computational powers. Furthermore, we conjecture that if the agents broadcast their beliefs rounded to a (fixed in advance) polynomial number of significant digits, then our techniques can be extended to establish a similar PSPACE-hardness result.

However, if one instead assumes that the beliefs are broadcast up to an arbitrary precision, our proof fails for a rather annoying reason: When implementing alternation from NP to Π_2 in the binary action model, if a formula ϕ has no satisfying assignments, we can exactly compute the belief of the NP observer agent. However, in case ϕ has a satisfying assignment, we can compute the belief only with high, but imperfect precision. The reason is that the exact value of the belief depends on the number of satisfying assignments of ϕ. This imperfection can be "rounded away" if the agents output a discrete guess for θ, but we do not know how to handle it if the beliefs are broadcast exactly.

Nevertheless, one can obtain a $\#\mathsf{P}$-hardness proof in the revealed belief model. The proof is by reduction from counting satisfying assignments in a 2-SAT formula. However, since the differences in belief corresponding to different numbers of satisfying assignments are small, it is not clear if they can be amplified, and consequently we do not demonstrate hardness of approximation (in that respect our result is similar to Papadimitriou and Tsitsiklis (1987)). For ease of exposition we introduce an additional relaxation to the model by allowing some agents to receive ternary private signals.

Theorem 15 *Assume the revealed belief model with beliefs transmitted up to arbitrary precision and call the respective computational problem* BINARY-BELIEF. *Additionally, assume that some agents receive ternary signals $\mathcal{S}(u) \in \{0, 1, 2\}$.*

There exists an efficient reduction that maps a 2-SAT *formula ϕ with N variables, M clauses and s satisfying assignments to an instance of* BINARY-BELIEF$(\Pi, t, u, H(u,t))$ *such that:*

- *The Bayesian network G has size $O(N + M)$, the time is set to $t = 2$ and agent u does not receive a private signal.*

- *All private signal probabilities come from a fixed family of at most ten distributions.*
- *The likelihood ratio of u at time $t = 2$ satisfies*

$$\frac{s}{2^N}\left(1-\frac{1}{4^N}\right) \le \frac{\mu(u,2)}{1-\mu(u,2)} \le \frac{s}{2^N}\left(1+\frac{1}{4^N}\right) .$$

In particular, rounding this ratio to the nearest multiple of 2^{-N} yields $s \cdot 2^{-N}$ and allows to recover s.

Reduction Our reduction uses the DAG structure and the concept of significant time as explained in Section A. The general idea is as in Theorem 1, with some adaptations to the counting setting and revealed beliefs. We assume that the agents broadcast beliefs in the form of LLRs.

We define a common signal distribution with $p_\mathsf{T} := 3/4$ and $p_\mathsf{F} := 1/4$ and respective LLRs ℓ_1 and ℓ_0. The graph we construct contains an observer agent OBS with no private signal and the "evaluation" agent EVAL with a $(p_\mathsf{T}, p_\mathsf{F})$ private signal. Given a 2-SAT formula ϕ with variables $x_1, \ldots, x_N$ and clauses $\eta_1, \ldots, \eta_M$, respective variable and clause gadgets are designed as follows:

For a variable x_i, we create two agents x_i and $\neg x_i$, receiving $(p_\mathsf{T}, p_\mathsf{F})$ private signals. Those two agents are observed by an auxiliary agent, which in turn is observed by agent OBS. The observation history of OBS indicates that the auxiliary agent broadcasts LLR equal to $\ell_0 + \ell_1$. At the same time, OBS observes another auxiliary agent with informative private signal, broadcasting LLR equal to $-\ell_0 - \ell_1$. See Figure 13 for illustration. Since the LLR broadcast by the agent observing x_i and $\neg x_i$ is the sum of their LLRs, we can perform an analysis similar to the threshold gadget in the binary action model. The result is that the variable gadget ensures that $\mathcal{S}(x_i) \neq \mathcal{S}(\neg x_i)$ and that each consistent signal configuration gives equal LLRs of $\theta = \mathsf{T}$ and $\theta = \mathsf{F}$.

Figure 13: Revealed belief reduction: Variable gadget.

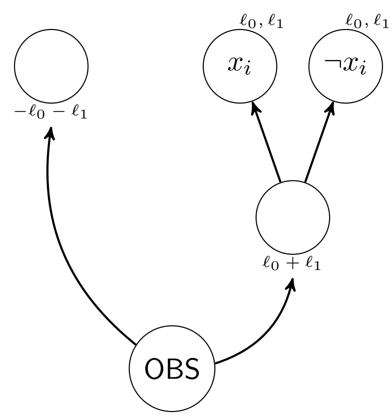

In the clause gadget (see Figure 14) for a clause η_j there is an auxiliary agent observing four agents:

- Two agents corresponding to the literals occurring in η_j.
- Agent EVAL.

- Agent F_j that receives a private signal $\mathcal{S}(F_j) \in \{0, 1, 2\}$. Its signal distribution is such that the respective LLRs satisfy

$$\begin{aligned} m_0 &:= \ell_0(F_j) = \ln \frac{\Pr[\mathcal{S}(F_j) = 0 \mid \theta = \mathsf{T}]}{\Pr[\mathcal{S}(F_j) = 0 \mid \theta = \mathsf{F}]} , \\ m_1 &:= \ell_1(F_j) = m_0 + \delta , \\ m_2 &:= \ell_2(F_j) = m_0 + 2\delta , \end{aligned} \tag{31}$$

where $\delta := \ell_1 - \ell_0 = 2 \ln 3$. Furthermore, the probabilities $q(\theta_0, b) := \Pr[\mathcal{S}(E_j) = b \mid \theta = \theta_0]$ for $\theta_0 \in \{\mathsf{T}, \mathsf{F}\}$ and $b \in \{0, 1, 2\}$ are chosen such that

$$q(\mathsf{T}, 2)q(\mathsf{T}, 0) = q(\mathsf{T}, 1)^2 = q(\mathsf{F}, 2)q(\mathsf{F}, 0) = q(\mathsf{F}, 1)^2 . \tag{32}$$

One checks that (31) and (32) are achieved (with $m_0 = -\delta$) by setting $q(\mathsf{T}, 2) = q(\mathsf{F}, 0) = q''$, $q(\mathsf{T}, 1) = q(\mathsf{F}, 1) = q'$, $q(\mathsf{T}, 0) = q(\mathsf{F}, 2) = 1 - q' - q''$, where (q'', q') is the unique positive solution of

$$\begin{cases} \frac{q''}{1-q'-q''} = 9 , \\ (q')^2 = q'(1 - q' - q'') , \end{cases}$$

which turns out to be $q'' = 9/13$ and $q' = 3/13$.

The auxiliary agent is observed by OBS, broadcasting belief $3\ell_0 + m_0 + 3\delta$.

Figure 14: Clause gadget.

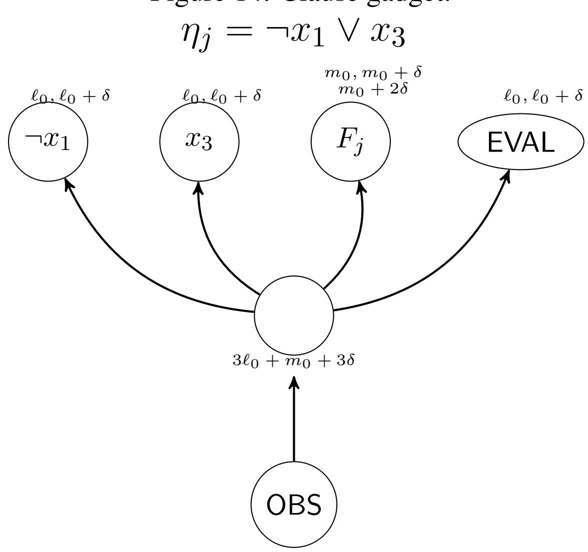

Since we want to be somewhat more precise in estimating LLRs induced by different assignments, we introduce additional gadgets "neutralizing" LLRs induced by signals of agents F_j and EVAL, illustrated in Figure 15. Their principle is basically the same as for the variable agents. For example, for each agent F_j we introduce another agent H_j with the same signal distribution, an agent observing both F_j and H_j and broadcasting $2m_0 + 2\delta$ to OBS and yet another agent broadcasting opposite belief $-2m_0 - 2\delta$ to OBS. In all, these agents ensure that any private signals to F_j and EVAL do not affect the LLR of the state of the world θ.

Figure 15: Gadgets for F_j and EVAL agents.

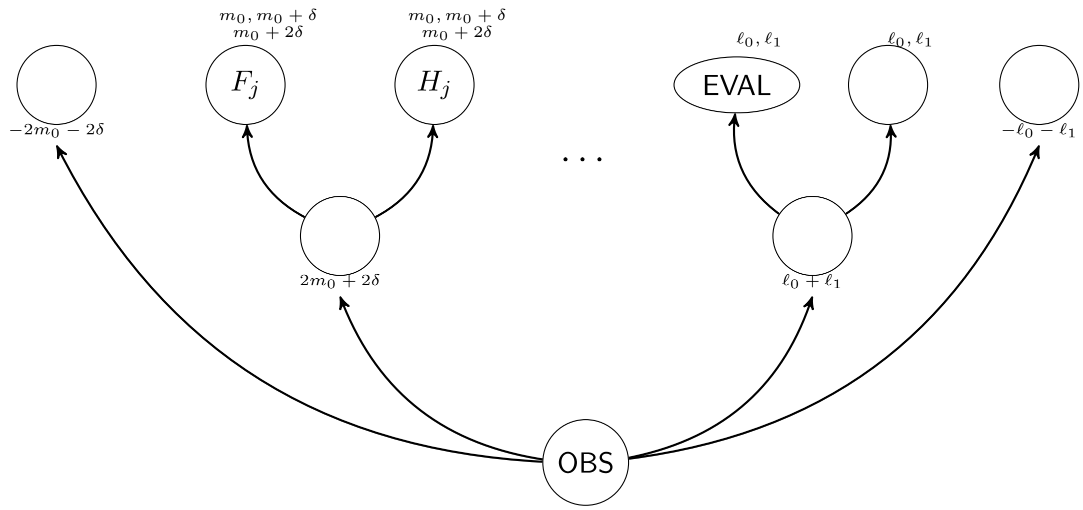

Finally, we let $r := 2N$ and introduce agents $B_1, \dots, B_r$ and $C_1, \dots, C_r$. Each agent B_i receives a $(p_\mathsf{T}, p_\mathsf{F})$ private signal. Agent C_i observes agents EVAL and B_i and broadcasts $\ell_0 + \ell_1$ to agent OBS (see Figure 16). This concludes the description of the reduction.

Figure 16: One of K parts of the "amplification" mechanism.

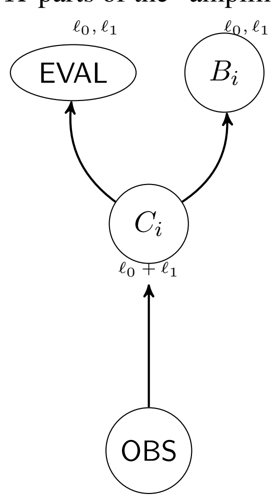

Analysis The analysis proceeds analogously to the proof of Theorem 1. First, the network is clearly of size $O(N + M)$ and has the required DAG structure with significant time $t = 2$ for agent OBS. Next, we convince ourselves that the private signals consistent with observations of OBS can be characterized as:

- For each assignment $\boldsymbol{v}$ there exists exactly one consistent configuration of private signals such that $\mathcal{S}(\mathsf{EVAL}) = 1$ and $\mathcal{S}(B_i) = 0$ for each $i \in \{1, \dots, r\}$.

- For each *satisfying* assignment $\boldsymbol{v}$ there is exactly one consistent configuration such that $\mathcal{S}(\mathsf{EVAL}) = 0$ and $\mathcal{S}(B_i) = 1$ for each $i \in \{1, \ldots, r\}$.
- There are no other consistent signal configurations.

Let us define $P(\boldsymbol{v}, b, \theta_0)$ as the probability that $\theta = \theta_0$ and that there arises the unique signal configuration consistent with assignment $\boldsymbol{v}$ and $\mathcal{S}(\mathsf{EVAL}) = b$. The gadgets (recall the relation (32) for agents F_j and H_j) ensure that $P(\cdot)$ is equal to

$$P(\boldsymbol{v}, 1, \mathsf{T}) = q \cdot \left(\frac{1}{4}\right)^r , \qquad P(\boldsymbol{v}, 1, \mathsf{F}) = q \cdot \left(\frac{3}{4}\right)^r ,$$

and, for each assignment x that is satisfying, additionally

$$P(\boldsymbol{v}, 0, \mathsf{T}) = q \cdot \left(\frac{3}{4}\right)^r , \qquad P(\boldsymbol{v}, 0, \mathsf{F}) = q \cdot \left(\frac{1}{4}\right)^r ,$$

where q is a universal common factor that depends only on N and M. Recalling that s denotes the number of satisfying assignments in ϕ, we can conclude that the LLR of agent OBS at its significant time $t = 2$ is given by

$$\begin{aligned} \frac{\mu(\mathsf{OBS})}{1 - \mu(\mathsf{OBS})} &= \frac{s \cdot (3/4)^r + 2^N \cdot (1/4)^r}{2^N \cdot (3/4)^r + s \cdot (1/4)^r} = \frac{s}{2^N} \cdot \frac{1 + \frac{2^N}{s \cdot 3^r}}{1 + \frac{s}{2^N \cdot 3^r}} \\ &\in \frac{s}{2^N} \cdot \left[1 - \frac{1}{3^r}, 1 + \frac{2^N}{3^r}\right] \subseteq \frac{s}{2^N} \cdot \left[1 \pm \frac{1}{4^N}\right] . \end{aligned}$$

In particular,

$$\left| \frac{\mu(\mathsf{OBS})}{1 - \mu(\mathsf{OBS})} - \frac{s}{2^N} \right| \le \frac{s}{8^N} < \frac{1}{2^{N+1}} ,$$

so rounding the likelihood ratio to the nearest multiple of 2^{-N} successfully recovers the number of satisfying assignments s. ■