

Actual Causality and Responsibility Attribution in Decentralized Partially Observable Markov Decision Processes

Stelios Triantafyllou
Max Planck Institute for
Software Systems
Germany
strianta@mpi-sws.org

Adish Singla
Max Planck Institute for
Software Systems
Germany
adishs@mpi-sws.org

Goran Radanovic
Max Planck Institute for
Software Systems
Germany
gradanovic@mpi-sws.org

ABSTRACT

Actual causality and a closely related concept of responsibility attribution are central to accountable decision making. Actual causality focuses on specific outcomes and aims to identify decisions (actions) that were critical in realizing an outcome of interest. Responsibility attribution is complementary and aims to identify the extent to which decision makers (agents) are responsible for this outcome. In this paper, we study these concepts under a widely used framework for multi-agent sequential decision making under uncertainty: decentralized partially observable Markov decision processes (Dec-POMDPs). Following recent works in RL that show correspondence between POMDPs and Structural Causal Models (SCMs), we first establish a connection between Dec-POMDPs and SCMs. This connection enables us to utilize a language for describing actual causality from prior work and study existing definitions of actual causality in Dec-POMDPs. Given that some of the well-known definitions may lead to counter-intuitive actual causes, we introduce a novel definition that more explicitly accounts for causal dependencies between agents' actions. We then turn to responsibility attribution based on actual causality, where we argue that in ascribing responsibility to an agent it is important to consider both the number of actual causes in which the agent participates, as well as its ability to manipulate its own degree of responsibility. Motivated by these arguments we introduce a family of responsibility attribution methods that extends prior work, while accounting for the aforementioned considerations. Finally, through a simulation-based experiment, we compare different definitions of actual causality and responsibility attribution methods. The empirical results demonstrate the qualitative difference between the considered definitions of actual causality and their impact on attributed responsibility.

CCS CONCEPTS

• **Computing methodologies** → **Multi-agent systems; Sequential decision making.**

KEYWORDS

Multi-Agent Systems, Actual Causality, Responsibility Attribution

ACM Reference Format:

Stelios Triantafyllou, Adish Singla, and Goran Radanovic. 2022. Actual Causality and Responsibility Attribution in Decentralized Partially Observable Markov Decision Processes. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society (AIES'22), August 1–3, 2022, Oxford, United Kingdom*. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3514094.3534133>

1 INTRODUCTION

Ex-post analysis of a decision making outcome, be it perceived positive or negative, is central to accountability, which is considered to be one of the pillars of trustworthy AI [11]. Such an analysis can enable us to pinpoint decisions (hereafter actions) that caused failures and assign responsibility to decision makers (hereafter agents) involved in the decision making process. When the emphasis is put on a specific outcome and circumstances, actions that were critical in realizing this outcome constitute *actual causes*. The extent to which the agents' actions were critical for the outcome of interest determines the agents' degrees of responsibility. Both actual causality and responsibility attribution have been well studied in moral philosophy, law, AI and related fields [3, 8–10, 17, 22, 27, 28, 32, 52].

A canonical approach to actual causality is based on the but-for test, which examines the counterfactual dependence of the outcome on agents' actions. It states that an action (or more generally, a set of actions) is a but-for cause of the outcome if the outcome would not have occurred had the action (resp. the set of actions) not been taken. It is well-known that but-for causes do not always align with human intuition—we refer the reader to [17] for an extensive discussion. Given this, much of the recent work on actual causality has tried to extend but-for causes in order to capture nuances of decision making scenarios where they seem to fail.

Some of the most influential extensions are due to Halpern and Pearl [16, 20, 21], who use Structural Causal Models (SCMs) [41] as a framework for reasoning about actual causality. Focusing on the modified Halpern-Pearl (HP) definition [16], actual causes are identified through an extended but-for test, evaluated relative to some contingency. As argued by Halpern [16], placing appropriate restrictions on contingencies is subtle; in the modified HP definition, contingencies can only be formed from non-causal variables set to their actual values.

While the modified HP definition generalizes but-for causality, it may still yield counter-intuitive actual causes when applied to sequential decision making [16]. An example that illustrates this is a variant of the *bogus prevention* scenario [27]. In this example, we have two agents, Assassin *A* and Bodyguard *B*, whose actions influence Victim *V*. By poisoning *V*'s coffee, *A* can cause *V*'s death, whereas *B* can prevent *V* from dying by putting an antidote. One



This work is licensed under a Creative Commons Attribution International 4.0 License.

AIES'22, August 1–3, 2022, Oxford, United Kingdom
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9247-1/22/08.
<https://doi.org/10.1145/3514094.3534133>

may ask, if A decides on its action after observing the action of B and only poisons V 's coffee if B has put the antidote, which actions should constitute the actual causes of V 's survival, in case B puts the antidote and then A poisons the coffee? As argued by Halpern [16], in contrast to what our intuitions would suggest, under the modified HP definition (as well as other variants of the HP definition), B 's action is an actual cause. Namely, putting the antidote passes the but-for test under the contingency that A poisons V 's coffee. However, this is not an answer that one would expect, since A had no intention of poisoning V in the first place. To correct for this, one may resort to normality considerations and extend the HP definition accordingly [18]. For example, when examining causality, the extended definition would exclude the “abnormal world” where B does not add the antidote and A poisons V 's coffee [16].

However, as we show in this paper, there are sequential decision making scenarios where the HP definitions provide counter-intuitive actual causes, even under the normality considerations. These novel scenarios demonstrate that existing definitions of actual causality (i.e., the but-for definition and the HP definitions) do not fully account for conditions under which an agent decides on its actions. These conditions generally depend on the interaction history, i.e., the previous actions of the agent or the other agents.

In this paper, we study actual causation in decentralized partially observable Markov decision processes (Dec-POMDPs) [38], which are widely used for modeling multi-agent interactions under uncertainty. Our goal is to utilize this framework in order to derive a novel definition of actual causality that more explicitly accounts for causal dependencies between agents' actions and their policies. As a down-stream task of interest, we consider responsibility attribution based on actual causality. Our contributions are as follows.

Framework. By relying on the recent results in reinforcement learning [7, 37], which show the correspondence between POMDPs and SCMs, we establish a connection between Dec-POMDPs and SCMs. This allows us to study existing definitions of actual causality and responsibility attribution methods in Dec-POMDPs.

Formal Properties. Using sequential decision making scenarios inspired by those from the moral philosophy literature, we argue that some of the most prominent definitions of actual causality (i.e., the but-for definition and the modified HP definition) do not fully account for causal dependencies between agents' actions. The corresponding nuances are formally captured by two novel properties: *Counterfactual Eligibility* and *Actual Cause-Witness Minimality*.

New Definition of Actual Causality. We then propose a definition of actual causality that satisfies the two novel properties. This definition utilizes additional variables, which are a part of the standard agent modeling approach in Dec-POMDPs [38] that assigns to each agent an information state specifying how the agent's policy depends on the interaction history.

Responsibility Attribution. We additionally study responsibility attribution based on actual causality. We introduce a family of responsibility attribution methods that extends the responsibility attribution method of Chockler and Halpern [8]. These methods take into consideration the number of actual causes an agent participates in and preserve a type of *performance incentive* akin to the one studied by [48]—an agent cannot reduce its own degree of responsibility by increasing the number of its actions that must be changed in order to obtain a different final outcome.

Experimental Results. Using a simulation-based experiment, we test the qualitative properties of different definitions of actual causation and we quantify their influence on responsibility assignments. The experimental results show that the modified HP definition violates the two novel properties rather frequently in one of the standard benchmarks for multi-agent RL—the card game *Goofspiel*. For example, for a game configuration in which agents can take 12 actions in total, we find that in the majority of trajectories, 2 or more actions (i.e., more than 16% of actions) do not conform to *Counterfactual Eligibility*. Similarly, 4 or more actual causes do not conform to *Actual Cause-Witness Minimality*. Note that the majority of trajectories have at least 13 actual causes. The but-for definition, which satisfies *Actual Cause-Witness Minimality*, violates *Counterfactual Eligibility* even more often than the HP definition: in the game configuration from above, 3 or more actions (i.e., 25%) do not conform to *Counterfactual Eligibility*. The results also show that these property violations can have a significant effect on agents' degrees of responsibility. When we correct for them, the agents' degrees of responsibility change in total by up to 50%-112%, depending on the responsibility attribution method.

We believe that these results shed a new light on actual causality and responsibility attribution, as they showcase additional challenges related to multi-agent sequential decision making. To the best of our knowledge, this is the first work that aims to tackle these challenges.

1.1 Related Work

In this subsection we provide a brief overview of the most relevant prior work, categorized in three different research topics: *actual causality*, *responsibility and blame attribution*, and *other works*.

Actual Causality. Arguably the closest to this paper is a recent line of work on actual causality in AI due to Halpern and Pearl [16, 20, 21], who introduced different versions of the HP definition of actual causality. Works that are closely related to the HP definitions are extensively surveyed in [16, 17], and they include: Pearl [40], who introduced the notion of *causal beam* that inspired the HP definitions; Hitchcock [26], who identifies a variable as an actual cause by searching for a causal path in which the variable passes the but-for test; Hall [14], who considers the *H-account* definition, which identifies a subset of actual causes identified by the HP definitions; and Halpern and Hitchcock [18] who extend the HP definitions by incorporating normality considerations. As we already mentioned, we extend this line of work by studying actual causality in Dec-POMDPs, which enables us to more explicitly model causal dependencies between agents' actions. This paper is also closely related to a more recent work by Baier et al. [2], who model multi-agent interaction via extensive form games, accounting for the conditions under which an agent decides on its actions through information states. However, Baier et al. [2] study orthogonal aspects, primarily focusing on responsibility attribution. In contrast, we contribute to the literature on actual causality by proposing a new definition that tackles challenges related to multi-agent sequential decision making, identified in this paper.

Responsibility and Blame Attribution. This paper is also related to the literature on responsibility attribution in multi-agent decision making. We already mentioned Chockler and Halpern [8],

who consider a causality-based notion of responsibility, and Baier et al. [3], who provide a game-theoretic account of the forward and backward notions of responsibility from [43]. Alechina et al. [1] extend the decision-oriented notion of responsibility from Chockler and Halpern [8] to assign responsibility to agents for the failure of a team plan [36, 51]. In our work, we use their method as a baseline for responsibility attribution. Yazdanpanah et al. [53] study a notion of responsibly akin to the notion of blame from [8],¹ and similar to this paper, they explicitly incorporate time. However, their framework is based on alternating-time temporal logic (ATL), whereas we utilize Dec-POMDPs, which are more suitable for decision making under uncertainty. Halpern and Kleiman-Weiner [19] formalize the notions of blameworthiness and intent using actual causality; similar to the degree of blame from [8] (see Footnote 1.1), these notions depend on the epistemic state of an agent. Friedenber and Halpern [12] extend the notion of blameworthiness to cooperative multi-agent settings. In contrast, this paper takes the notion of responsibility defined by Chockler and Halpern [8] as its starting point. The work by Triantafyllou et al. [48] is perhaps the closest in spirit to this paper as it studies blame attribution in multi-agent Markov decision processes. However, the focus of that work is on average performance as an outcome of interest, whereas we focus on specific events along a decision making trajectory. Naturally, this paper broadly relates to (cooperative) game theory and cost sharing games [30, 50], since attribution methods such as Shapley value [46, 47] or Banzhaf index [4, 5] are often utilized for defining degrees of blame, responsibility and blameworthiness [3, 12, 48].

Other Works. From a technical point of view, this paper closely relates to RL approaches that utilize SCMs. Buesing et al. [7] leverages SCMs for policy evaluation, which in turn can improve policy search methods in model-based RL. Oberst and Sontag [37] extend the framework of Buesing et al. [7], allowing for off-policy evaluation in POMDPs with stochastic transition dynamics. Madumal et al. [34] utilize causal models to generate explanations for actions taken by a RL agent. Tsirtsis et al. [49] consider a causal model of the environment based on Markov decision processes (MDPs). They use this model to find an alternative sequence of actions that maximizes the counterfactual outcome, but is within a certain Hamming distance from the original action sequence. This alternative sequence serves as a counterfactual explanation. We contribute to this line of work by establishing a connection between Dec-POMDPs and SCMs and utilizing it for actual causality and responsibility attribution in multi-agent sequential decision making. Finally, this paper relates to the recent work on counterfactual *credit assignment* in RL [23, 35], where the goal is to improve an agent’s learning efficiency by properly crediting an action for its effect on the obtained rewards. Our focus is not on improving the learning process of an agent, but on accountability considerations.

2 FORMAL SETTING AND PRELIMINARIES

In this section, we describe our formal setting, based on decentralized partially observable Markov decision processes (Dec-POMDPs) [6, 38] and structural causal models (SCMs) [41, 42]. We also review

and adopt to our setting a language for reasoning about actual causality [17], and we formally model the actual causality problem in the context of multi-agent sequential decision making.

2.1 Decentralized Partially Observable Markov Decision Processes (Dec-POMDPs)

We consider a Dec-POMDP $\mathcal{M} = (\mathcal{S}, \{1, \dots, n\}, \mathcal{A}, \mathcal{P}, O, \Omega, T, \sigma)$ with n agents, where: \mathcal{S} is the state space; $\{1, \dots, n\}$ is the agents’ set; $\mathcal{A} = \times_{i=1}^n \mathcal{A}_i$ is the joint action space, with \mathcal{A}_i being the action space of agent i ; P specifies transitions with $P(s, a, s')$ denoting the probability of the process transitioning to s' from s when agents $\{1, \dots, n\}$ take joint action $a = (a_1, \dots, a_n)$; $O = \times_{i=1}^n O_i$ is the joint observation space, with O_i being the observation space of agent i ; Ω is an observation probability function with $\Omega(s, o)$ denoting the probability of agents $\{1, \dots, n\}$ receiving joint observation $o = (o_1, \dots, o_n)$ when in state s ; T is the finite time horizon; σ is the initial state distribution. We assume \mathcal{S} , \mathcal{A} and O to be finite and discrete. For ease of notation, we additionally assume that the agents’ immediate rewards are part of their observations. Throughout the paper, we denote random variables with capital letters, e.g., S , A and O .

We also consider for each agent i a model $m_i = (I_i, \pi_i, Z_i, Z_{i,0})$ [38], where: I_i is the (finite and discrete) information state space of i ; π_i is the policy of agent i , i.e., a mapping $\pi_i : I_i \rightarrow \Delta(\mathcal{A}_i)$, where $\Delta(\mathcal{A}_i)$ is a probability simplex over \mathcal{A}_i ; Z_i is agent i ’s information probability function with $Z_i(i_t, a_i, o_i, i'_t)$ denoting the probability of i ’s information state changing from i_t to i'_t , after i takes action a_i and observes o_i ; $Z_{i,0}$ is i ’s initial information probability function depending only on $o_{i,0}$. We use $\pi_i(a_i|i_t)$ to denote the probability of agent i taking action a_i given information state i_t . The agents’ joint policy is denoted by π , and we assume that $\pi(a|l_1, \dots, l_n) = \pi_1(a_1|l_1) \cdot \dots \cdot \pi_n(a_n|l_n)$. Note that information states are a way to encode the information that an agent uses in its decision making.

2.2 Dec-POMDPs and Structural Causal Models

Although Dec-POMDPs are a very general and useful modelling tool for multi-agent sequential decision making, they are not sufficient to reason counterfactually about alternate outcomes [33], and hence actual causality.² For instance, given a trajectory $\tau = \{(s_t, a_{1,t}, \dots, a_{n,t})\}_{t=0}^{T-1}$ generated by Dec-POMDP \mathcal{M} under joint policy π , we would like to predict what would have happened, had agent i taken action a'_i instead of action $a_{i,t}$. However, even though we have access through \mathcal{M} to the probability distribution of the next state $P(s_t, (a_{1,t}, \dots, a'_{i,t}, \dots, a_{n,t}), \cdot)$, we do not have a way to infer what would be the value of S_{t+1} .³ Following Buesing et al. [7], to overcome this limitation we view \mathcal{M} under joint policy π as a structural causal model (SCM) C . To do this, we express P , Ω , $\{Z_i\}_{i \in \{1, \dots, n\}}$ and $\{\pi_i\}_{i \in \{1, \dots, n\}}$ as deterministic functions g with independent noise variables U , such as

$$\begin{aligned} S_t &= g_{S_t}(S_{t-1}, A_{t-1}, U_{S_t}), & O_t &= g_{O_t}(S_t, U_{O_t}), \\ I_{i,t} &= g_{I_{i,t}}(I_{i,t-1}, A_{i,t-1}, O_{i,t}, U_{I_{i,t}}), & A_{i,t} &= g_{A_{i,t}}(I_{i,t}, U_{A_{i,t}}), \end{aligned} \quad (1)$$

where U_{S_t} , U_{O_t} , $U_{I_{i,t}}$ and $U_{A_{i,t}}$ are $|S|$ -, $|O|$ -, $|I_i|$ - and $|A_i|$ - dimensional, respectively. It can be shown that such a parameterization

¹Chockler and Halpern [8] differentiate responsibility and blame. For example, one of the key difference is that an agent’s degree of blame depends on its epistemic state (i.e., the agent’s belief about the underlying causal model).

²By counterfactual reasoning, we mean predicting what would have happened in a specific instance of the decision process (trajectory) had some action(s) been different.
³Or in general the value of anything that comes (chronologically) after time-step t .

is always possible.⁴ Henceforth, we will refer to SCMs that are defined in this way as Dec-POMDP SCMs. Consistent with the SCMs' terminology [41], we also say that state variables S_t , observation variables O_t , information variables $I_{i,t}$ and action variables $A_{i,t}$ constitute the endogenous variables of C , U are the exogenous variables, and equations (1) are the model's structural equations. The causal graph of the Dec-POMDP SCM can be found in Appendix A. We can generate a trajectory $\tau = \{(s_t, a_{1,t}, \dots, a_{n,t})\}_{t=0}^{T-1}$ using Dec-POMDP SCM by simply specifying a setting \vec{u} for the exogenous variables in U , also called context, and then solving the structural equations of C , i.e., Eq. (1). Note that for each Dec-POMDP SCM-context pair (C, \vec{u}) , also called causal setting, there is a unique trajectory τ that can be generated in that way. Importantly, we can also find out what would have happened in τ , had agent i taken action a'_i instead of $a_{i,t}$ in the following way:

- (1) We perform the intervention⁵ $A_{i,t} \leftarrow a'_i$ on C , that is we replace $g_{A_{i,t}}(I_{i,t}, U_{A_{i,t}})$ in Eq. (1) with the value a'_i . The resulting SCM is denoted by $C^{A_{i,t} \leftarrow a'_i}$.
- (2) We generate the counterfactual trajectory τ^{cf} from the causal setting $(C^{A_{i,t} \leftarrow a'_i}, \vec{u})$, where \vec{u} is the same context that we used to generate τ .⁶

Note that when Dec-POMDP \mathcal{M} or joint policy π are stochastic, the counterfactual trajectory τ^{cf} may not be identifiable without further assumptions [37]. This is because, there may be multiple parameterizations of a Dec-POMDP SCM, i.e., multiple functions g and distributions over the exogenous variables U , which are all able to correctly represent \mathcal{M} under π ,⁷ but which suggest different counterfactual outcomes, e.g., τ^{cf} . Consequentially, the choice of model can have a significant impact on claims of causality. In our experiments, we choose to focus on a particular class of SCMs, the Gumbel-Max SCMs [37]. This class of SCMs has been shown to satisfy the desirable property of counterfactual stability, which excludes a specific type of non-intuitive counterfactual outcomes. Appendix B provides more details on Gumbel-Max SCMs and the counterfactual stability property.

2.3 Actual Causality

We now review and adopt in our formal setting a language introduced by prior work on actual causality with SCMs [17]. Consider a Dec-POMDP SCM C and a context \vec{u} , and the (unique) trajectory generated by the causal setting (C, \vec{u}) , $\tau = \{(s_t, a_t)\}_{t=0}^{T-1}$. We call primitive event, any formula of the form $V = v$, where V is an endogenous variable in C , i.e., state, observation, information or action variable, and v is a valid value for V . Let ϕ be an event, that is any Boolean combination of primitive events. We use $(C, \vec{u}) \models \phi$ to denote that ϕ is true in the causal setting (C, \vec{u}) , i.e., ϕ takes place in τ . Furthermore, given a set of interventions on action variables $\vec{A} \leftarrow \vec{a}'$, we write $(C, \vec{u}) \models [\vec{A} \leftarrow \vec{a}']\phi$, if $(C^{\vec{A} \leftarrow \vec{a}'}, \vec{u}) \models \phi$. For instance, consider the counterfactual scenario in which agent i

takes the action a'_i instead of $a_{i,t}$ in τ . If under this counterfactual scenario the process transitions to state s at time-step $t + 1$, then the following statement holds

$$(C, \vec{u}) \models [A_{i,t} \leftarrow a'_i](S_{t+1} = s).$$

Actual Causality in Multi-Agent Sequential Decision Making. Our goal is to pinpoint the actions which caused a particular event to happen. Given a causal setting (C, \vec{u}) and the event of interest ϕ , we want to determine the actual causes of ϕ in (C, \vec{u}) . In this paper, what can be an actual cause is a conjunction of primitive events consisted only of action variables, abbreviated here as $\vec{A} = \vec{a}$. We say that every conjunct of actual cause $\vec{A} = \vec{a}$ is part of that cause. Furthermore, in some cases we want to define an actual cause $\vec{A} = \vec{a}$ w.r.t. some contingency, that is $\vec{A} = \vec{a}$ is an actual cause only if that contingency holds. What can be a contingency in this paper is again a conjunction of primitive events consisted only of action variables, abbreviated as $\vec{W} = \vec{w}'$. Finally, for $\vec{A} = \vec{a}$ to be an actual cause of ϕ in (C, \vec{u}) under contingency $\vec{W} = \vec{w}'$, there has to exist a setting \vec{a}' , such that $(C, \vec{u}) \models [\vec{A} \leftarrow \vec{a}', \vec{W} \leftarrow \vec{w}']\neg\phi$. We will often refer to \vec{a}' as (counterfactual) setting. Consistent with the actual causality literature [17], we call the tuple $(\vec{W}, \vec{w}', \vec{a}')$ a witness to the fact that $\vec{A} = \vec{a}$ is an actual cause of ϕ in (C, \vec{u}) .

Coming back to the introduction example, (C, \vec{u}) models the considered trajectory: B puts the antidote ($B = \textit{antidote}$); A poisons V 's coffee ($A = \textit{poison}$); V survives. The outcome of interest ϕ is that V survives. According to the HP definition, the action $B = \textit{antidote}$ consists the actual cause of ϕ in (C, \vec{u}) under the contingency that $A = \textit{poison}$. Indeed, there is a counterfactual setting for B such that

$$(C, \vec{u}) \models [B = \textit{not antidote}, A = \textit{poison}]\neg(V \textit{ survives}).$$

In Section 3, we consider several definitions of actual causality w.r.t. a causal setting (C, \vec{u}) , where C is always a Dec-POMDP SCM. More specifically, in this paper an actual causality definition \mathcal{D} has to formally describe a process that receives as input a causal setting (C, \vec{u}) and an event of interest ϕ , and outputs a set of actual cause-witness pairs, i.e., a set of elements of the form $(\vec{A} = \vec{a}, (\vec{W}, \vec{w}', \vec{a}'))$. We use $H_{\mathcal{D}}((C, \vec{u}), \phi)$ (or just $H_{\mathcal{D}}$, when (C, \vec{u}) and ϕ are implied) to denote the set of all actual cause-witness pairs of ϕ in (C, \vec{u}) , according to \mathcal{D} . We also refer to (C, \vec{u}) as the actual world or situation. Similarly, we refer to a causal setting as the counterfactual world or scenario when it results from (C, \vec{u}) after an intervention is performed on a subset of its action variables, e.g., $(C^{\vec{A} \leftarrow \vec{a}'}, \vec{u})$.

3 DEFINITIONS FOR ACTUAL CAUSE

In this section, we analyze two of the most popular definitions of actual causality that involve counterfactuals, the “but-for” definition⁸ (BF definition from now on) [22], and the Halpern and Pearl definition (HP definition from now on) [16]. We provide two counterexamples (both are new variants of the “bogus prevention” scenario [25]) which expose several weaknesses of the two definitions. We formally capture the insights we gain from these examples with two novel properties. Subsection 3.3 introduces a new definition for actual cause, which satisfies these two properties.

⁸Also known as cause-in-fact and *sine qua non*.

⁴In [7], they show how to represent an episodic POMDP as an SCM, and prove that this is always possible. Their results can be directly extended to Dec-POMDPs.

⁵Interventions are also often modeled through the *do*-operator [39], $do(A_{i,t} = a'_i)$.

⁶Definitions for actual causality and responsibility, which are the focus of this paper, are relative to a causal setting [8]. Therefore, we assume C and \vec{u} to be fully known.

⁷For every state, observation, information or action variable V , it holds that variable V' equals V in distribution, where $V' = g_V(pa_V, U_V)$, and pa_V are the parents of V in the causal graph.

3.1 The BF Definition

But-for cause is one of the fundamental definitions of causation in law [22], and it states that C is a cause of E if but for C , E would not have occurred. In other words, C was necessary for E to happen. In our setting, we formally define but-for cause as follows.

Definition 3.1. (But-For Cause) $\vec{A} = \vec{a}$ is a but-for cause of the event ϕ in (C, \vec{u}) if the following conditions hold:

BFC1. $(C, \vec{u}) \models (\vec{A} = \vec{a})$ and $(C, \vec{u}) \models \phi$

BFC2. There is a setting \vec{a}' of the variables in \vec{A} , such that

$$(C, \vec{u}) \models [\vec{A} \leftarrow \vec{a}'] \neg \phi$$

BFC3. \vec{A} is minimal; there is no subset \vec{A}' of \vec{A} , such that $\vec{A}' = \vec{a}'$ satisfies BFC1 and BFC2, where \vec{a}' is the restriction of \vec{a} to the variables of \vec{A}

We say that $(\emptyset, \emptyset, \vec{a}')$ is a witness of $\vec{A} = \vec{a}$ being a but-for cause of ϕ in (C, \vec{u}) .

BFC1 requires that both $\vec{A} = \vec{a}$ and ϕ happened in the actual world, (C, \vec{u}) . BFC2 implies the but-for condition, i.e., but for $\vec{A} = \vec{a}$, ϕ would not have occurred. BFC3 is a minimality condition, which ensures that an actual cause does not include any non-essential elements. Unfortunately, but-for cause does not suffice for a good definition of actual cause in the context of sequential decision making, and the next example illustrates some of the reasons.

Example 3.2. Victim V dines at time-step t . Assassin A has access to V 's table at time-steps $t - 2$ and $t - 1$, when they can choose whether to *poison* or *not poison* V 's water. A 's policy is to always *poison* V 's water, unless it is already poisoned. We consider the trajectory, in which A chooses to *poison* V 's water only at time-step $t - 2$, and V dies from the poison at time t .

To identify a but-for cause of V dying at time-step t , consider an intervention that sets A 's action at time-step $t - 2$ to *not poison*. If this is the only intervention, A follows its policy at $t - 1$ and takes action *poison*, which results in the same outcome. To change the outcome, we also need to intervene at time-step $t - 1$ and set A 's action to *not poison*. This implies that action *poison* taken at $t - 2$ and action *not poison* taken at $t - 1$ form a but-for cause of V dying at time t . We find this counter-intuitive because the action that has to be changed at time-step $t - 1$ is not the one that was taken in the actual situation but the one that would have been taken in the counterfactual scenario where A does *not poison* the water at $t - 2$. Since the conditions that influence A 's decision at $t - 1$ change once we intervene at $t - 2$, we argue that the action taken at $t - 1$ should not be a part of an actual cause, but should be treated as a contingency. The following property formalizes this insight.

Property 3.3 (Counterfactual Eligibility). We say that a definition for actual cause \mathcal{D} satisfies *Counterfactual Eligibility* if for every (C, \vec{u}) , ϕ , $(\vec{A} = \vec{a}, (\vec{W}, \vec{w}', \vec{a}'))$ and $A_{i,t} = a_{i,t}$, where $(\vec{A} = \vec{a}, (\vec{W}, \vec{w}', \vec{a}'))$ is an actual cause-witness pair of ϕ in (C, \vec{u}) according to \mathcal{D} , and $A_{i,t} = a_{i,t}$ is part of $\vec{A} = \vec{a}$, it holds that $(C, \vec{u}) \models [\vec{A} \leftarrow \vec{a}', \vec{W} \leftarrow \vec{w}'](I_{i,t} = i_{i,t})$, where $i_{i,t}$ is agent i 's information state in (C, \vec{u}) , i.e., $(C, \vec{u}) \models (I_{i,t} = i_{i,t})$.

Property 3.3 states that an agent's action is eligible for being a part of an actual cause if the information state under which the

agent took that action in the actual world remains the same in the witness world, i.e., the counterfactual world which corresponds to the cause's witness. As Example 3.2 suggests, the BF definition violates Property 3.3.

3.2 The HP Definition

Arguably, one of the most influential accounts of causality is Halpern and Pearl's notion of actual causes in SCMs [21]. There are three variants of the HP definition of actual causality [16, 20, 21]. In this paper we consider and adopt in our setting the latest one [16].⁹

Definition 3.4. (HP) $\vec{A} = \vec{a}$ is an actual cause of the event ϕ in (C, \vec{u}) if the following conditions hold:

HPC1. $(C, \vec{u}) \models (\vec{A} = \vec{a})$ and $(C, \vec{u}) \models \phi$

HPC2. There is a set \vec{W} of action variables and a setting \vec{a}' of the variables in \vec{A} such that if $(C, \vec{u}) \models (\vec{W} = \vec{w}')$, then

$$(C, \vec{u}) \models [\vec{A} \leftarrow \vec{a}', \vec{W} \leftarrow \vec{w}'] \neg \phi$$

HPC3. \vec{A} is minimal; there is no subset \vec{A}' of \vec{A} , such that $\vec{A}' = \vec{a}'$ satisfies HPC1 and HPC2, where \vec{a}' is the restriction of \vec{a} to the variables of \vec{A}

We say that $(\vec{W}, \vec{w}', \vec{a}')$ is a witness of $\vec{A} = \vec{a}$ being an actual cause of ϕ in (C, \vec{u}) .

HPC1 and HPC3 are similar to BFC1 and BFC3, respectively. HPC2 says that the but-for condition holds under the contingency $\vec{W} = \vec{w}'$, where the setting \vec{w}' has the observed value of \vec{W} in (C, \vec{u}) . Roughly speaking, this means that $\vec{A} = \vec{a}$ is an actual cause of ϕ in (C, \vec{u}) if but for $\vec{A} = \vec{a}$, ϕ would not have happened, had the action variables in \vec{W} been fixed to their actual values. The main intuition behind HPC2, and what differentiates this HP definition from its predecessors, is that “*only what happens in the actual situation should matter*” [17].

Coming back to Example 3.2, according to the HP definition, the actual cause of V dying at time-step t is the action of A to *poison* the water at time-step $t - 2$, under the contingency that A would *not poison* the water at $t - 1$. In other words, if we assume that in the counterfactual world where A does *not poison* the water at $t - 2$, they also do *not poison* the water at $t - 1$, then the first action is considered an actual cause of V dying. We find this answer more intuitive than the one given by the BF definition. Despite the success of the HP definition in Example 3.2 as well as in many more examples in the moral philosophy literature [16, 17], we illustrate with the next example that the HP definition is not sufficient for multi-agent sequential decision making.

Example 3.5. Victim V dines at time-step t . Bodyguard B , who suspects a poisonous attack, has access to V 's table at time-step $t - 2$, when they can choose where to put an antidote, either into V 's water or into V 's wine. B is right, indeed an assassin A has access to V 's table at time-step $t - 1$, when they can choose where to put the poison, again either into V 's water or into V 's wine. The poison is neutralized by the antidote only if they have been put into the same liquid, otherwise V dies. We assume that A observes where B puts the antidote and that its intention is to poison V . Therefore, A 's policy is to put the poison into the liquid that does not have the antidote. Consider the trajectory, in which B puts the antidote

⁹All definitions and their relations are extensively discussed in [17].

into the *water* at time-step $t - 2$, A puts the poison into the *wine* at time-step $t - 1$, and V dies at time t .

According to both BF and HP definitions, an actual cause of V dying at t is the action of A poisoning the *wine* at $t - 1$. However, according to the HP definition, this is not the only actual cause of V 's death. The action of B putting the antidote into the *water* is also considered an actual cause, under the contingency that A would poison the *wine*, i.e., B 's action is an actual cause assuming that A 's action is fixed to its value in the "actual situation".¹⁰ In particular, B did put the antidote into the *water* in the actual scenario, and since it is a single action, it is also minimal, hence *HPC1* and *HPC3* are satisfied. Furthermore, if we intervene on B 's action by changing it to *wine*, and fix A 's action also to *wine* (the realized action), then V does not die, and hence *HPC2* is satisfied.

We find the latter actual cause to be counter-intuitive. Essentially B had no control over the final outcome of this example, because of the full observability and policy assumed for A . This counter-intuitive result is due to the fact that the HP definition applies the minimality condition (*HPC3*) only on the variables of the actual cause $\vec{A} = \vec{a}$, and does not include those of the contingency $\vec{W} = \vec{w}'$ (here $\vec{W} = \vec{w}'$ is in fact another actual cause). Motivated by Example 3.5, we introduce the following property.

Property 3.6 (Actual Cause-Witness Minimality). We say that a definition for actual cause \mathcal{D} satisfies *Actual Cause-Witness Minimality* if for every (C, \vec{u}) , ϕ , $\vec{A} = \vec{a}$ and $\vec{W} = \vec{w}'$, where $\vec{A} = \vec{a}$ is an actual cause of ϕ in (C, \vec{u}) under contingency $\vec{W} = \vec{w}'$ according to \mathcal{D} , there are no $\vec{A}' = \vec{a}'$ and $\vec{W}' = \vec{w}''$, such that $\vec{A}' \cup \vec{W}' \subset \vec{A} \cup \vec{W}$ and $\vec{A}' = \vec{a}'$ is an actual cause of ϕ in (C, \vec{u}) under contingency $\vec{W}' = \vec{w}''$ according to \mathcal{D} , where \vec{a}' is the restriction of \vec{a} to \vec{A}' .

Roughly speaking, Property 3.6 extends *HPC3* to also include \vec{W} , i.e., $\vec{A} \cup \vec{W}$ is minimal. As Example 3.5 suggests, the HP definition violates Property 3.6. Additionally, Appendix C.1 describes a scenario where the HP definition violates Property 3.3.

Example 3.5 also sets the ground for arguing about the interpretation of the *actuality* test: "only what happens in the actual situation should matter". For instance, one can argue that A 's action of putting the poison into the *wine* is a valid contingency for the HP definition since this action did realize in the actual situation. Arguably, this interpretation is adopted in [16]. In contrast, we argue for an interpretation that focuses not just on agents' actions, but also on their information states, i.e., conditions under which agents reach their decisions. Under this interpretation, A 's action of putting the poison into the *wine* does not pass the actuality test for the counterfactual world in which B puts the antidote into the *wine*. Namely, in the actual situation, A put the poison into the *wine* only because B had put the antidote into the *water*. Now, this interpretation may be restrictive if the actuality test is applied on contingencies, as it is the case with the HP definition. For instance, if condition *HPC2* is modified accordingly, the HP definition would identify no actual causes in Example 3.2. However, we believe that the actuality test is not important for contingencies, but only for actual causes. This is formalized by Property 3.3. Note also that the BF definition satisfies Property 3.6 because of *BFC3*.

¹⁰Note that based on Theorem 2.3 from [16] both older versions of the HP definition also consider B 's action as an actual cause of V 's death in Example 3.5.

Normality and Defaults. The notions of normality and defaults have been shown to deal with a number of examples where the HP definitions provide counter-intuitive actual causes [15, 18]. However, Example 3.5 is not one of them. More specifically, in this example there are two possible worlds, one where B and A put the antidote and the poison into the same liquid, and one where they put them into different liquids. Given the intentions of A and that A observes B 's action in this example, one may consider the former world less normal than the latter one. All 3 HP definitions, when extended to account for the aforementioned normality ordering, they provide no actual causes, although there is arguably one. We conclude, that despite the usefulness of these notions, they do not address the shortcomings of the core definition described above.

3.3 A New Definition for Actual Cause

We extend the BF definition with the notion of contingencies and implement the insights we gain from Examples 3.2 and 3.5, to propose a new definition for actual cause. Intuitively, our definition takes a but-for cause and splits its set of (action) variables into two subsets: the actual cause and the contingency. The partition is based on whether the conditions under which these actions were taken, change between the actual world and the witness one.

Definition 3.7. (Actual Cause) $\vec{A} = \vec{a}$ is an actual cause of the event ϕ in (C, \vec{u}) , under the contingency $\vec{W} = \vec{w}'$ if the following conditions hold:

- AC1. There is a setting \vec{a}' of the variables in \vec{A} , such that $\vec{A} = \vec{a} \wedge \vec{W} = \vec{w}$ is a but-for cause of ϕ in (C, \vec{u}) , and also satisfies condition *BFC2* with setting (\vec{a}', \vec{w}')
- AC2. For every agent i and time-step t such that $A_{i,t} \in \vec{A}$ and $(C, \vec{u}) \models (I_{i,t} = i_{i,t})$, it holds that

$$(C, \vec{u}) \models [\vec{A} \leftarrow \vec{a}', \vec{W} \leftarrow \vec{w}'] (I_{i,t} = i_{i,t})$$
- AC3. For every agent i and time-step t such that $A_{i,t} \in \vec{W}$ and $(C, \vec{u}) \models (I_{i,t} = i_{i,t})$, it holds that

$$(C, \vec{u}) \models [\vec{A} \leftarrow \vec{a}', \vec{W} \leftarrow \vec{w}'] \neg (I_{i,t} = i_{i,t})$$

We say that $(\vec{W}, \vec{w}', \vec{a}')$ is a witness of $\vec{A} = \vec{a}$ being an actual cause of ϕ in (C, \vec{u}) .

AC1 says that combined $\vec{A} = \vec{a}$ and $\vec{W} = \vec{w}$ should form a but-for cause of ϕ in (C, \vec{u}) using the settings \vec{a}' and \vec{w}' . According to the BF definition, this means that the following conditions should hold:

- (1) $(C, \vec{u}) \models (\vec{A} = \vec{a})$, $(C, \vec{u}) \models (\vec{W} = \vec{w})$ and $(C, \vec{u}) \models \phi$
- (2) $(C, \vec{u}) \models [\vec{A} \leftarrow \vec{a}', \vec{W} \leftarrow \vec{w}'] \neg \phi$
- (3) $\vec{A} \cup \vec{W}$ is minimal w.r.t. conditions (1) and (2)

AC2 requires that the actual cause $\vec{A} = \vec{a}$ should contain only (action) variables for which their underlying conditions (information states) in the counterfactual world $(C^{\vec{A} \leftarrow \vec{a}', \vec{W} \leftarrow \vec{w}'}, \vec{u})$ are the same as in the actual world (C, \vec{u}) . *AC3* says that the contingency $\vec{W} = \vec{w}'$ should contain only variables for which these conditions change. Regarding Example 3.2, our definition agrees with the HP definition, because of condition *AC3*. Regarding Example 3.5, our definition agrees with the BF definition because of condition *AC1*. Furthermore, conditions *AC2* and *AC1* guarantee Properties 3.3 and 3.6, respectively.

Definition 3.7 can also be extended with the notion of normality, by following the same procedure as with the HP definition in [18].

4 RESPONSIBILITY

In this section, we study approaches to determining the agents' degree of responsibility relative to a causal setting (C, \vec{u}) and an event of interest ϕ . We focus on approaches that assign responsibility based on the actual causes of ϕ in (C, \vec{u}) . More specifically, in Section 4.1, we adopt in our setting and analyze a well-known definition of responsibility introduced by Chockler and Halpern [8]. In Section 4.2, we introduce a new family of definitions that extend the Chockler and Halpern definition.

4.1 The Chockler-Halpern Definition

Chockler and Halpern [8] show that the original HP definition of causality [20] can be used to assign a degree of responsibility to each primitive event, measuring how pivotal it was for the event of interest. Halpern [17] modifies the Chockler-Halpern definition (hereafter CH) to incorporate the other two HP definitions [16, 21]. Alechina et al. [1] extend the analysis by appraising the responsibility of each agent. In our framework and for a definition of causality \mathcal{D} , the CH notion of responsibility can be defined as follows.

Definition 4.1. (CH) Consider a causal setting (C, \vec{u}) and an event of interest ϕ such that $(C, \vec{u}) \models \phi$. Agent i 's degree of responsibility for ϕ in (C, \vec{u}) is 0 if none of i 's actions is a part of an actual cause according to \mathcal{D} . Otherwise it is the maximum value $\frac{m}{k}$ such that if $\vec{A} = \vec{a}$ is an actual cause of ϕ in (C, \vec{u}) under the contingency $\vec{W} = \vec{w}'$ according to \mathcal{D} , then $k = |\vec{A}| + |\vec{W}|$ and m denotes the number of agent i 's action variables in \vec{A} .

The CH definition captures the important idea that an agent's degree of responsibility should depend on the size of the actual causes it participates in, the size of their corresponding contingencies, and its degree of participation. However, as mentioned by Baier et al. [2], the CH definition does not take into consideration the number of actual causes an agent is involved in, which is evidence of that agent's power over the final outcome. Additionally, the definition also ignores other aspects of actual causality that one might consider important for attributing responsibility, such as the number of different contingencies an actual cause might have.

4.2 A Family of Methods that Extend CH

We consider the CH definition and extend it in a natural way, so that an agent's degree of responsibility is now determined by a wider variety of actual causes, instead of just one. More specifically, our new definition takes into account the whole set of actual cause-witness pairs $H_{\mathcal{D}}$ for some definition \mathcal{D} and applies weight vectors over that set. These vectors are non-negative and agent-specific, and they determine by how much an agent's degree of responsibility is affected by each pair in $H_{\mathcal{D}}$. Each weight vector \vec{b} has to have at least one strictly positive element.

Definition 4.2. (Degree of Responsibility) Consider a causal setting (C, \vec{u}) and an event of interest ϕ such that $(C, \vec{u}) \models \phi$. Given a weight vector \vec{b} over the set $H_{\mathcal{D}}$, agent i 's degree of responsibility for ϕ in (C, \vec{u}) is 0 if none of i 's actions is part of an actual cause

according to \mathcal{D} ; otherwise it is

$$\frac{\sum_{c \in \{1, \dots, |H_{\mathcal{D}}|\}} b_c \cdot \frac{m_c}{k_c}}{\sum_{c \in \{1, \dots, |H_{\mathcal{D}}|\}} b_c}$$

such that if $(\vec{A} = \vec{a}, (\vec{W}, \vec{w}', \vec{a}'))$ is the c -th actual cause-witness pair of $H_{\mathcal{D}}$, then $k_c = |\vec{A}| + |\vec{W}| - w_c$, and m_c and w_c denote the number of agent i 's action variables in \vec{A} and \vec{W} , respectively.

Definition 4.2 is flexible in the sense that it can generate different responsibility attribution methods by changing the agents' weight vectors. Compared to Definition 4.1, an agent's degree of responsibility does not depend anymore on the number of action variables the agent has in a contingency of an actual cause in which it participates. In simpler words, our definition guarantees that an agent would not be attributed reduced responsibility had it adopted a policy that would make more "mistakes" in the counterfactual scenario.¹¹ For instance, in Example 3.2, if A 's policy was to *poison* K 's water only at time-step $t - 2$ then its degree of responsibility according to CH (and \mathcal{D} being either the HP definition or Definition 3.7) would be 1. However, it would be 1/2 if its policy was to always *poison* the water. On the contrary, for responsibility attribution methods from Definition 4.2, the agent's degree of responsibility is 1 in both cases. Note that, similar to CH, an agent's degree of responsibility according to Definition 4.2 is always between 0 and 1.¹² More specifically, if the agent had no impact on the outcome, its degree of responsibility would be 0, while if it was the only agent with full control over the outcome, its responsibility would be 1.

5 EXPERIMENTS

In this section, we experimentally test the studied definitions of actual causality (Section 3) and responsibility attribution methods (Section 4) using an experimental testbed based on the card game *Goofspiel*. Appendix E contains additional experimental results.

5.1 Environment and Policies

The game. *Goofspiel*(N) is a two-person card game where each player's initial hand consists of the cards $\{1, \dots, N\}$. There is a face down central pile of cards (also $\{1, \dots, N\}$) called the deck, which is shuffled in the beginning of each game. In every round, the top card of the deck is flipped. Then, both players choose a card from their hand and simultaneously reveal it. The player with the higher card wins the round, and in the case of a tie no player wins. If a player manages to win the round, they are awarded a number of points equal to the value of the flipped card, also called the prize, otherwise they are awarded 0 points. All cards played in that round are then discarded and a new round starts. After N rounds, the player with the most points wins the game. Note that typically $N = 13$, making the mathematical analysis of the game quite challenging [44, 45]. Moreover, it is worth mentioning that *Goofspiel*(N) is part of a well known framework for RL in games [24, 31]. We introduce a version of this game which has two teams of two players. We call this version *TeamGoofspiel*(N). The game proceeds as before, with the difference that now the team which cumulatively bids the higher cards in a round is the team that obtains that round's prize.

¹¹This guarantee is aligned with the intuition behind the blame attribution property of *Performance Monotonicity*, introduced by Triantafyllou et al. [48].

¹²The agents' degrees of responsibility do not have to sum up to 1 [17].

The players. We assume the agency over the members of one of the teams, whose players are referred to as agents and are denoted by Ag_0 and Ag_1 . We treat the other team as a part of the environment, and we refer to its members as opponents. All the players are assumed to keep similar information states at each time-step/round. More specifically, the information state of a player at each round consists of: the remaining cards on their hand; the round’s prize; partial information about the current score—if their team is winning or not. Notice that for simplicity, in this setting players do not keep track of which cards the other players discarded in previous rounds, i.e., they don’t condition their actions on the available moves of other players, nor they try to infer their policies.

Policies. The policy of Ag_0 is to always match the round’s prize whenever possible. We differentiate two cases when this action is not feasible, i.e., if the matching card is not on Ag_0 ’s hand. If their team is winning (resp. not winning) they play the card with the highest (resp. lowest) value out of the cards with a value lower (resp. higher) than the prize. In case such a card is not available, they play the card with the lowest (resp. highest) value on their hand. The policy of Ag_1 is to always play the card with the highest value on their hand, if the prize is greater than their hand’s average card value minus X , otherwise they play the card with the lowest value on their hand. X is 0 if their team is winning and 1 otherwise.

Both opponents follow the same stochastic policy defined as follows. If their team is winning (resp. not winning) they randomly choose a card from their hand with value lower (resp. higher) or equal to the prize, and if they don’t have such a card in their hand they randomly choose one of the cards available.

We choose the players’ policies to follow simple rules, and depend on small size information states, so that the generated actual causes are easy to interpret, but not too simple and small, so that they become trivial. Note also that the random and the matching policies are quite standard in *Goofspiel(N)* analysis, and that the latter have been shown to dominate the former [13, 45].

Actual Causes. We focus on trajectories, i.e., instances of the game in which the final outcome is either a win for the opponents’ team or a draw. For each of these trajectories, our goal is to pinpoint those actions of the agents that caused them to not win the game, and then quantify the agents’ influence on that outcome. In particular, we specify the set of all actual cause-witness pairs for each trajectory, and based on this set we compute the agents’ degrees of responsibility. Note that in order to generate a trajectory in the *TeamGoofspiel(N)* environment, we first need to sample from the initial state distribution, i.e., shuffle the deck, and then at each time-step sample the opponents’ actions based on the distributions defined by their stochastic policies.

5.2 Demonstration Example

In this section, we focus on a particular trajectory of *TeamGoofspiel(N)* and present the set of actual cause-witness pairs for that trajectory based on Definition 3.7. To compute this set, we implement a simple tree search algorithm that iterates over all possible alternative actions of the agents. In our experiments, we restrict the size of actual cause-witness pairs to 4, in order to obtain more interpretable actual causes. Namely, large actual causes suggest a counterfactual world that is quite different from what actually

Table 1: Actual Cause-Witness Pairs Based on Definition 3.7

Actual Cause	CF Setting	Contingency	Improvement
$A_{0,1} = 4, A_{1,1} = 4$	1, 1	-	2
$A_{0,1} = 4, A_{1,3} = 2$	1, 1	$A_{0,2} = 2$	1
	2, 1	$A_{0,2} = 1$	1
$A_{0,1} = 4, A_{1,2} = 3$	1, 1	-	2
$A_{1,1} = 4, A_{0,3} = 2$	1, 1	-	1
		-	1
$A_{1,1} = 4, A_{0,2} = 3$	1, 1	$A_{1,2} = 3$	2
		$A_{1,2} = 4$	1
$A_{0,0} = 5$	1	-	1
	2	-	1
$A_{1,0} = 5$	1	$A_{0,3} = 1$	1
		$A_{1,3} = 4$	1
		$A_{1,1} = 3$	3
	2	$A_{1,1} = 4$	1
		$A_{1,2} = 3$	1
		$A_{1,1} = 3$	3
		$A_{1,1} = 4$	1
3	$A_{1,2} = 2$	1	
$A_{1,0} = 5, A_{0,2} = 3$	1, 2	-	1
$A_{1,0} = 5, A_{0,1} = 4$	1, 2	-	1

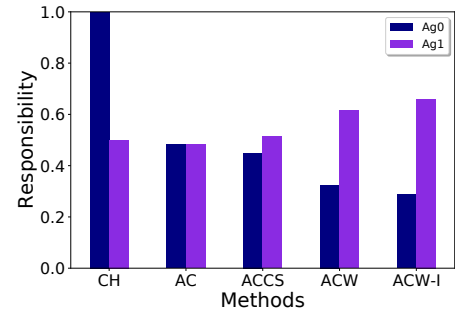


Figure 1: Responsibility Attribution Based on Table 1

happened, and they are also more difficult to comprehend. Thus, smaller actual causes are usually preferred.

For $N = 5$, we consider the trajectory where prizes show up in a descending order (5, 4, 3, 2, 1). Both agents’ and opponents’ actions in this scenario always match the prize, resulting in a 0 – 0 draw. Table 1 shows the set of all actual cause-witness pairs for that trajectory. Interestingly, and despite its simplicity, the considered trajectory admits 19 different actual cause-witness pairs, each of them representing a set of minimal changes that the agents could have made in order to win the game. More specifically, each row of Table 1 corresponds to one actual cause-witness pair $(\vec{A} = \vec{a}, (\vec{W}, \vec{w}', \vec{a}'))$, where: column 1 includes the actual causes $\vec{A} = \vec{a}$; column 2 includes the counterfactual settings \vec{a}' ; column 3 includes the contingencies $\vec{W} = \vec{w}'$; column 4 includes the improvement in score difference the agents achieve in the corresponding counterfactual worlds.

The corresponding tables for definitions BF and HP can be found in Appendix C.1.

5.3 Instances of Definition 4.2

As mentioned in Section 4.2, by changing the values of the agent-specific weight vectors \vec{b} in Definition 4.2 we can generate multiple

different responsibility attribution methods. Next, we present the instances of Definition 4.2 we consider in our experiments. The names of the corresponding attribution methods are derived from the elements of actual causality which they consider distinct.

AC: For every actual cause $\vec{A} = \vec{a}$, the weight of exactly one pair ($\vec{A} = \vec{a}, (\vec{W}, \vec{w}', \vec{a}')$) is 1, and all other weights are 0. For Table 1, this means that we keep one row per actual cause, and we discard every other row for that actual cause and the column Improvement. For an agent i , the actual cause-witness pair we choose for each actual cause, i.e., the pair whose weight is 1, is one with the highest value of $\frac{m_c}{k_c}$ for that agent. This method takes into account the number of distinct actual causes an agent participates in.

ACCS: For every actual cause $\vec{A} = \vec{a}$ and counterfactual setting \vec{a}' , the weight of exactly one pair ($\vec{A} = \vec{a}, (\vec{W}, \vec{w}', \vec{a}')$) is 1, and all other weights are 0. For Table 1, this means that we keep exactly one row per actual cause-counterfactual setting pair, and we discard every other row for that pair and the column Improvement. For an agent i , the actual cause-witness pair we choose for each actual cause-counterfactual setting pair, i.e., the pair whose weight is 1, is one with the highest value of $\frac{m_c}{k_c}$ for that agent. Additional to **AC**, **ACCS** takes also into account the number of all the different counterfactual actions that the agents who participate in $\vec{A} = \vec{a}$ could have taken in order for the final outcome to improve.

The remaining attribution methods assume the same weight vectors for all agents.

ACW: The weight of each actual cause-witness pair is 1. For Table 1, this means that we take its whole content into account except for the column Improvement. Additional to **ACCS**, **ACW** takes into account all the different contingencies under which $\vec{A} = \vec{a}$ is an actual cause.

ACW-I: The weight of each actual cause-witness pair is equal to the value of the counterfactual improvement it admits. This means that we use the full information given in this table.¹³

Apart from **AC**, **ACCS**, **ACW** and **ACW-I**, we also consider in our experiments the CH definition. Plot 1 shows the agents' degrees of responsibility for the trajectory from Section 5.2 and for the various responsibility attribution methods. For this plot, the input of all the methods is Table 1. Observe how in this example the lion's share of responsibility shifts gradually from Ag0 to Ag1, as we include more information from Table 1 to our responsibility assignment process. For instance, Ag1 could improve the outcome on their own by playing one of 3 alternative actions at the first time-step (rows 10-16), while Ag0 had only 2 (rows 8, 9). Because of that, Ag1's responsibility increases relative to Ag0's when we transition from **AC** to **ACCS**. Appendix C.2 shows the attributed responsibilities, when BF and HP are considered instead of Definition 3.7.

5.4 Violations of Properties 3.3 and 3.6

In this section, we compute the frequency of Property 3.3 and Property 3.6 violations by the BF and HP definitions from Section 3. Furthermore, we examine by how much these property violations might affect the agents' degrees of responsibility. We measure both quantities for $N \in \{4, 5, 6, 7, 8\}$, and 50 trajectories per value of N .

¹³This approach to responsibility is aligned with the notion of graded causality [18]. However, here instead of using a normality ranking over the actual cause-witness pairs, we evaluate them w.r.t. the counterfactual improvement they admit.

5.4.1 Property 3.3 Violations. Plot 2a summarizes the frequency results for Property 3.3. More specifically, for each trajectory we compute the number of actions that are, according to Property 3.3, wrongfully characterized as part of one or more actual causes by the BF and HP definitions. For instance, consider the boxplot which corresponds to $N = 6$ and HP. For half of the trajectories, the HP definition considers at least 2 out of the total 12 actions as part of one or more actual causes, when it should had instead considered them as part of their contingencies.

Next, we want to measure by how much this mislabeling of actions, i.e., Property 3.3 violations, can affect the process of responsibility attribution. In order to quantify this measure, we execute the following procedure. For each trajectory, we first compute the set of actual cause-witness pairs $H_{\mathcal{D}}$ based on definition \mathcal{D} , where \mathcal{D} can be either BF or HP. Then, for every approach from Section 5.3 we compute the agents' degrees of responsibility utilizing the set $H_{\mathcal{D}}$. Next, we correct $H_{\mathcal{D}}$ for Property 3.3 violations, i.e., for every actual cause-witness pair ($\vec{A} = \vec{a}, (\vec{W}, \vec{w}', \vec{a}')$) in $H_{\mathcal{D}}$ we remove from $\vec{A} = \vec{a}$ and \vec{a}' all actions that violate Property 3.3, and we add them to the contingency set $\vec{W} = \vec{w}'$. We then take the newly defined set of actual cause-witness pairs $H'_{\mathcal{D}}$, and recompute the agents' degrees of responsibility. Plot 2b shows the maximum value of the total absolute difference between the two computed degrees of responsibility, for every value of N and responsibility method. The maximum is taken over all trajectories. We choose to plot the maximum differences to showcase the potential magnitude of unfairness that violating Property 3.3 might cause to the responsibility assignment. The results demonstrate that correcting BF and HP for these violations can have a significant impact on the agents' degrees of responsibility. It is also worth noting that the CH definition seems to be the least resilient to this type of violations among the definitions we consider here.

5.4.2 Property 3.6 Violations. Apart from Property 3.3, the HP definition also violates Property 3.6 (Section 3.2). Plot 2c displays the frequency of these violations (brown boxplots). More specifically, the plot shows for all trajectories the number of distinct actual cause-contingency pairs which are non-minimal, according to Property 3.6. While comparing this number to the total number of these pairs which is shown in Plot 2d, we conclude that the HP definition systematically violates Property 3.6.

To measure the impact of Property 3.6 violations on responsibility attribution, we follow a procedure similar to the one for Property 3.3 in Section 5.4.1. More specifically, we first compute the agents' degrees of responsibility based on $H_{\mathcal{D}}$, where \mathcal{D} is the HP definition. Next, we correct $H_{\mathcal{D}}$ for Property 3.6 violations, i.e., we remove all actual cause-witness pairs that violate Property 3.6, and we recompute the degrees of responsibility. Similar to Plot 2b, Plot 2e shows the maximum value of the total absolute difference between the two computed degrees of responsibility, for every value of N and responsibility method. The maximum is taken over all trajectories. The results indicate that Property 3.6 violations in the HP definition can greatly affect the downstream task of responsibility attribution. However, it is worth mentioning that for CH, the agents' degrees of responsibility changed only for 1 trajectory out of the 250 we sampled in our experiments, indicating that it is the most resilient to this type of property violations.

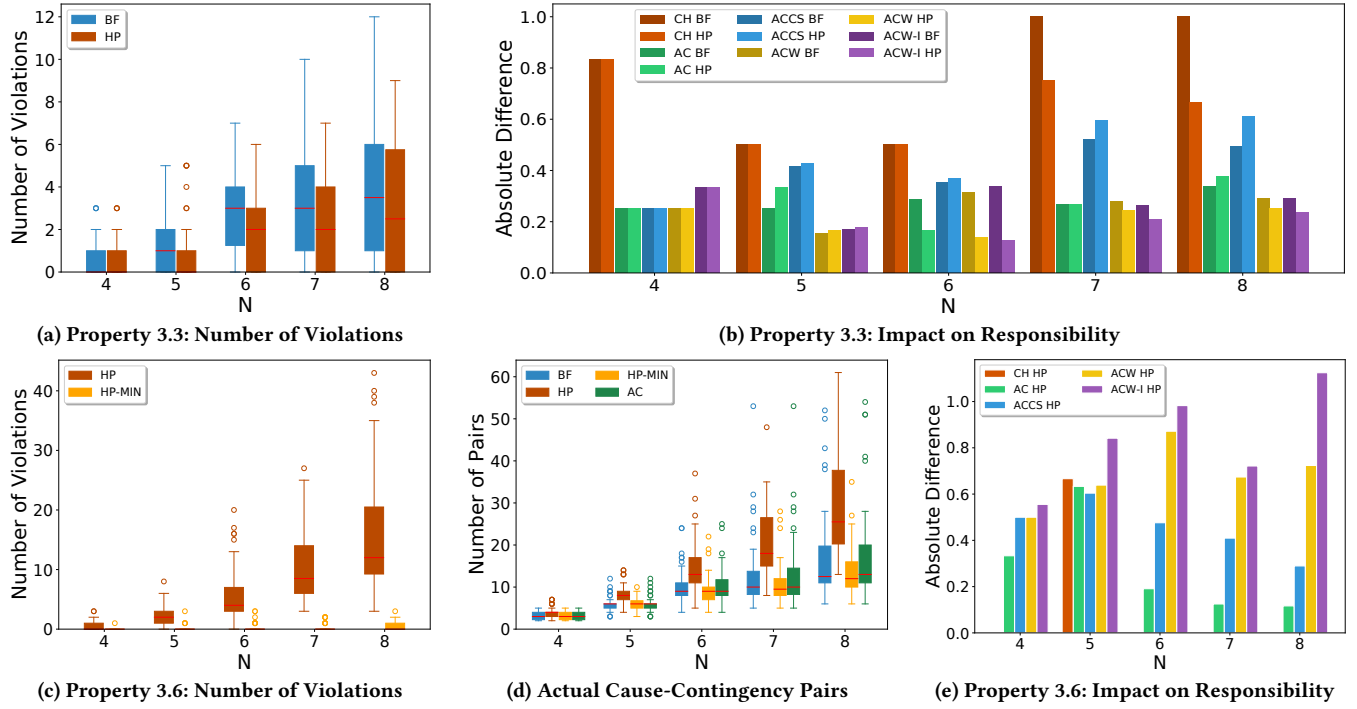


Figure 2: Plots 2a and 2c show the number of violations of Properties 3.3 and 3.6. Plots 2b and 2e show the impact of these violations on the agents’ degrees of responsibility. Plot 2d shows the number of distinct actual cause-contingency pairs.

Note that the HP definition allows for non-minimal contingencies, that is $(\vec{A} = \vec{a}, (\vec{W}, \vec{w}', \vec{a}'))$ may be considered as a valid actual cause-witness pair by HP even if there is a subset \vec{W}' of \vec{W} and a setting \vec{w}'' such that $(\vec{A} = \vec{a}, (\vec{W}', \vec{w}'', \vec{a}'))$ is also an actual cause-witness pair according to HP. As mentioned by Ibrahim [29], when attributing responsibility based on the HP definition it would make sense to impose the witness minimality condition in addition to *HPC3*. We denote this enhanced version of the HP definition by HP-MIN. Note that violations of the contingency minimality condition fall under Property 3.6 violations. Therefore, we expect that HP-MIN will do better than HP w.r.t Property 3.6 violations, and hence have a lower impact on responsibility.¹⁴ Plots 2c and 2d verify this intuition. They show that the number of violations is significantly smaller for HP-MIN. However, these violations are not completely vanished, meaning that there are still cases where they can have a large impact on responsibility attribution. In Appendix D, we plot again 2a, 2b and 2e, after replacing HP with HP-MIN.

6 CONCLUSION

To summarize, in this paper we studied actual causality and responsibility attribution through the lens of sequential decision making in Dec-POMDs. We identified some of the shortcomings of existing definitions of actual causality and introduced a new definition to address them. Furthermore, we extended one of the well known causality-based approaches to responsibility attribution in order to account for an agent’s power over the final outcome and its

ability to manipulate its own degree of responsibility. While this work focuses on particular challenges in defining actual causality and attributing responsibility, we view it as an important step toward establishing a formal framework that supports accountable multi-agent sequential decision making.

Some of the most interesting future research directions are related to practical considerations. Given that our primary goal is to formalize the notions of actual causality and responsibility attribution, we made simplifying assumptions that allowed more efficient computation of experimental results. For example, even though the Dec-POMDP framework adopted in this work does model uncertainty, we assumed the full knowledge of random variables that define contexts of the corresponding SCM. We also assumed that the agents’ policies are given. Lifting these assumptions is critical for making this work more applicable in practice. Furthermore, the algorithmic solution for inferring actual causes and assign responsibility in the experiments is based on exhaustive search. Therefore, deriving more scalable algorithmic solutions is needed for applying this work in challenging domains. Finally, we deem further analysis of actual causality properties a meaningful extension of our work.

ACKNOWLEDGEMENTS

This research was, in part, funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – project number 467367360. We thank the anonymous reviewers for their valuable comments and suggestions.

¹⁴In particular, only ACCS, ACW and ACW-I are affected.

REFERENCES

- [1] Natasha Alechina, Joseph Y. Halpern, and Brian Logan. 2020. Causality, responsibility and blame in team plans. *arXiv preprint arXiv:2005.10297* (2020).
- [2] Christel Baier, Florian Funke, and Rupak Majumdar. 2021. A game-theoretic account of responsibility allocation. *arXiv preprint arXiv:2105.09129* (2021).
- [3] Christel Baier, Florian Funke, and Rupak Majumdar. 2021. Responsibility attribution in parameterized Markovian models. In *Proc. of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 11734–11743.
- [4] John F. Banzhaf III. 1964. Weighted voting doesn't work: A mathematical analysis. *Rutgers Law Review* 19 (1964), 317.
- [5] John F. Banzhaf III. 1968. One man, 3.312 votes: a mathematical analysis of the Electoral College. *Villanova Law Review* 13 (1968), 304.
- [6] Daniel S. Bernstein, Robert Givan, Neil Immerman, and Shlomo Zilberstein. 2002. The complexity of decentralized control of Markov decision processes. *Mathematics of operations research* 27, 4 (2002), 819–840.
- [7] Lars Buesing, Theophane Weber, Yori Zwols, Sebastien Racaniere, Arthur Guez, Jean-Baptiste Lespiau, and Nicolas Heess. 2018. Woulda, coulda, shoulda: Counterfactually-guided policy search. *arXiv preprint arXiv:1811.06272* (2018).
- [8] Hana Chockler and Joseph Y. Halpern. 2004. Responsibility and blame: A structural-model approach. *Journal of Artificial Intelligence Research* 22 (2004), 93–115.
- [9] Mark Coeckelbergh. 2020. Artificial intelligence, responsibility attribution, and a relational justification of explainability. *Science and engineering ethics* 26, 4 (2020), 2051–2068.
- [10] Anupam Datta, Deepak Garg, Dilsun Kaynar, Divya Sharma, and Arunesh Sinha. 2015. Program actions as actual causes: A building block for accountability. In *2015 IEEE 28th Computer Security Foundations Symposium*. 261–275.
- [11] European Commission. 2019. Ethics Guidelines for Trustworthy Artificial Intelligence. URL: <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai>. [Online; accessed 15-January-2021].
- [12] Meir Friedenberg and Joseph Y. Halpern. 2019. Blameworthiness in multi-agent settings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 525–532.
- [13] Mark Grimes and Moshe Dror. 2013. Observations on strategies for Goofspiel. In *2013 IEEE Conference on Computational Intelligence in Games (CIG)*. 1–2.
- [14] Ned Hall. 2007. Structural equations and causation. *Philosophical Studies* 132, 1 (2007), 109–136.
- [15] Joseph Y. Halpern. 2008. Defaults and normality in causal structures. In *Eleventh International Conference of Knowledge Representation and Reasoning*. 198–208.
- [16] Joseph Y. Halpern. 2015. A modification of the Halpern-Pearl definition of causality. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*. 3022–3033.
- [17] Joseph Y. Halpern. 2016. *Actual causality*. MIT Press.
- [18] Joseph Y. Halpern and Christopher Hitchcock. 2015. Graded causation and defaults. *The British Journal for the Philosophy of Science* 66, 2 (2015), 413–457.
- [19] Joseph Y. Halpern and Max Kleiman-Weiner. 2018. Towards formal definitions of blameworthiness, intention, and moral responsibility. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [20] Joseph Y. Halpern and Judea Pearl. 2001. Causes and explanations: A structural-model approach. Part I: Causes.. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*. 194–202.
- [21] Joseph Y. Halpern and Judea Pearl. 2005. Causes and explanations: A structural-model approach. Part I: Causes. *British Journal for the Philosophy of Science* 56, 4 (2005).
- [22] Herbert Lionel Adolphus Hart and Tony Honoré. 1985. *Causation in the Law*. OUP Oxford.
- [23] Anna Harutyunyan, Will Dabney, Thomas Mesnard, Mohammad Gheshlaghi Azar, Bilal Piot, Nicolas Heess, Hado P. van Hasselt, Gregory Wayne, Satinder Singh, Doina Precup, and et.al. 2019. Hindsight credit assignment. *Advances in Neural Information Processing Systems* 32 (2019), 12467–12476.
- [24] Daniel Hennes, Dustin Morrill, Shayegan Omidshafiei, Rémi Munos, Julien Pérolat, Marc Lanctot, Audrunas Gruslys, Jean-Baptiste Lespiau, Paavo Parmas, Edgar Duéñez-Guzmán, et al. 2020. Neural replicator dynamics: Multiagent learning via hedging policy gradients. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*. 492–501.
- [25] Eric Hiddleston. 2005. Causal powers. *The British journal for the philosophy of science* 56, 1 (2005), 27–59.
- [26] Christopher Hitchcock. 2001. The intransitivity of causation revealed in equations and graphs. *The Journal of Philosophy* 98, 6 (2001), 273–299.
- [27] Christopher Hitchcock. 2007. Prevention, preemption, and the principle of sufficient reason. *The Philosophical Review* 116, 4 (2007), 495–532.
- [28] David Hume. 2000. *An enquiry concerning human understanding: A critical edition*. Oxford University Press.
- [29] Amjad Ibrahim. 2021. *An actual causality framework for accountable systems*. Ph. D. Dissertation. Technische Universität München.
- [30] Kamal Jain and Mohammad Mahdian. 2007. Cost sharing. *Algorithmic Game Theory* 15 (2007), 385–410.
- [31] Marc Lanctot, Edward Lockhart, Jean-Baptiste Lespiau, Vinicius Zambaldi, Satyaki Upadhyay, Julien Pérolat, Sriram Srinivasan, Finbarr Timbers, Karl Tuyls, Shayegan Omidshafiei, et al. 2019. OpenSpiel: A framework for reinforcement learning in games. *arXiv preprint arXiv:1908.09453* (2019).
- [32] David Lewis. 1974. Causation. *The Journal of Philosophy* 70, 17 (1974), 556–567.
- [33] David Lewis. 2013. *Counterfactuals*. John Wiley & Sons.
- [34] Prashan Madumal, Tim Miller, Liz Sonenberg, and Frank Vetere. 2020. Explainable reinforcement learning through a causal lens. In *Proceedings of the AAAI conference on artificial intelligence*. 2493–2500.
- [35] Thomas Mesnard, Theophane Weber, Fabio Viola, Shantanu Thakoor, Alaa Saade, Anna Harutyunyan, Will Dabney, Thomas S Stepleton, Nicolas Heess, Arthur Guez, and et.al. 2021. Counterfactual credit assignment in model-free reinforcement learning. In *International Conference on Machine Learning*. 7654–7664.
- [36] Roberto Micalizio, Pietro Torasso, and Gianluca Torta. 2004. On-line monitoring and diagnosis of multi-agent systems: A model based approach. In *Proceedings of the 16th European Conference on Artificial Intelligence*, Vol. 16. 848.
- [37] Michael Oberst and David Sontag. 2019. Counterfactual off-policy evaluation with gumbel-max structural causal models. In *International Conference on Machine Learning*. 4881–4890.
- [38] Frans A. Oliehoek and Christopher Amato. 2016. *A concise introduction to decentralized POMDPs*. Springer.
- [39] Judea Pearl. 1995. Causal diagrams for empirical research. *Biometrika* 82, 4 (1995), 669–688.
- [40] Judea Pearl. 1998. *On the definition of actual cause*. Technical Report R-259. Department of Computer Science, University of California, Los Angeles.
- [41] Judea Pearl. 2009. *Causality*. Cambridge University Press.
- [42] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. 2017. *Elements of causal inference: Foundations and learning algorithms*. The MIT Press.
- [43] Ibo van de Poel. 2011. The relation between forward-looking and backward-looking responsibility. In *Moral responsibility*. Springer, 37–52.
- [44] Glenn C. Rhoads and Laurent Bartholdi. 2012. Computer solution to the game of pure strategy. *Games* 3, 4 (2012), 150–156.
- [45] Sheldon M. Ross. 1971. Goofspiel—the game of pure strategy. *Journal of Applied Probability* 8, 3 (1971), 621–625.
- [46] Lloyd S. Shapley. 2016. *17. A value for n-person games*. Princeton University Press.
- [47] Lloyd S. Shapley and Martin Shubik. 1954. A method for evaluating the distribution of power in a committee system. *The American Political Science Review* 48, 3 (1954), 787–792.
- [48] Stelios Triantafyllou, Adish Singla, and Goran Radanovic. 2021. On blame attribution for accountable multi-agent sequential decision making. *Advances in Neural Information Processing Systems* 34 (2021).
- [49] Stratis Tsirtsis, Abir De, and Manuel Rodriguez. 2021. Counterfactual explanations in sequential decision making under uncertainty. *Advances in Neural Information Processing Systems* 34 (2021).
- [50] John Von Neumann and Oskar Morgenstern. 2007. *Theory of games and economic behavior (commemorative edition)*. Princeton University Press.
- [51] Cees Witteveen, Nico Roos, Roman van der Krogt, and Mathijs de Weerd. 2005. Diagnosis of single and multi-agent plans. In *Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*. 805–812.
- [52] Richard W. Wright. 1985. Causation in tort law. *Calif. L. Rev.* 73 (1985), 1735.
- [53] Vahid Yazdanpanah, Mehdi Dastani, Natasha Alechina, Brian Logan, and Wojciech Jamroga. 2019. Strategic responsibility under imperfect information. In *Proceedings of the 18th International Conference on Autonomous Agents and Multi-agent Systems AAMAS 2019*. 592–600.

A CAUSAL GRAPH OF DEC-POMDP SCM

In this section, we provide the causal graph of the Dec-POMDP SCM from Section 2.2. The graph is shown in Figure 3.

B GUMBEL-MAX SCMS AND COUNTERFACTUAL STABILITY

In this section, we show how Gumbel-Max SCMs can be implemented in our framework, and also provide the main intuition behind the counterfactual stability property. For more details on Gumbel-Max SCMs and the counterfactual stability property we refer the interested reader to [37].

Under the Gumbel-Max model, the structural equations of Eq. (1) become:

$$\begin{aligned} S_t &= \arg \max_{s \in \mathcal{S}} \{\log P(S_{t-1}, A_{t-1}, S_t = s) + U_{S_t}\} \\ O_t &= \arg \max_{o \in \mathcal{O}} \{\log \Omega(S_t, O_t = o) + U_{O_t}\} \\ I_{i,t} &= \arg \max_{i_i \in \mathcal{I}_i} \{\log Z_i(I_{i,t-1}, A_{i,t-1}, O_t, I_{i,t} = i_i) + U_{I_{i,t}}\} \\ A_{i,t} &= \arg \max_{a_i \in \mathcal{A}_i} \{\log \pi_i(A_{i,t} = a_i | I_{i,t}) + U_{A_{i,t}}\} \end{aligned}$$

where $U_{S_t}, U_{O_t}, U_{I_{i,t}}$ and $U_{A_{i,t}} \sim \text{Gumbel}(0, 1)$.

The class of Gumbel-Max SCMs has been shown to satisfy the desirable property of counterfactual stability, which excludes a specific type of non-intuitive counterfactual outcomes. We provide the main intuition behind this property, with the help of an example. Consider the observed trajectory $\tau = \{(s_t, a_t)\}_{t=0}^{T-1}$, and the counterfactual scenario in which agents $\{1, \dots, n\}$ take the joint action a' at time-step t , instead of a_t . The counterfactual stability property ensures that under this counterfactual scenario, it is impossible that at time-step $t+1$ the process would transition to a state s' different than the observed state, i.e., s_{t+1} if

$$\frac{P(S_t = s, A_t = a', S_{t+1} = s_{t+1})}{P(S_t = s, A_t = a_t, S_{t+1} = s_{t+1})} \geq \frac{P(S_t = s, A_t = a', S_{t+1} = s')}{P(S_t = s, A_t = a_t, S_{t+1} = s')}.$$

In other words, in order for the state at time-step $t+1$ to change under a counterfactual scenario, the relative likelihood of an alternative state s' must have increased compared to that of s_{t+1} .

C RESULTS OF SECTIONS 5.2 AND 5.3 FOR THE BF AND HP DEFINITIONS

C.1 Actual Cause-Witness Pairs

In this section, we provide the sets of actual cause-witness pairs for the trajectory considered in Section 5.2, and definitions BF (Table 2) and HP (Table 3).

We also describe a scenario where the HP definition violates Property 3.3. The actual cause-witness pair of Table 3 denoted by red color fails to meet the conditions stated by Property 3.3. More specifically, Ag0 's information state at time-step 2 is different between the actual world and the witness world, i.e., the agent's hand in the actual situation at time 2 is (1, 2, 3) and in the counterfactual scenario is (2, 3, 4). Despite that, action $A_{0,2} = 4$ is still considered as a part of the actual cause by the HP definition. Therefore, this scenario shows that the HP definition does not satisfy Property 3.3.

C.2 Responsibility

In this section, we provide the agents' degrees of responsibility for the trajectory considered in Section 5.2, and definitions BF (Plot 4) and HP (Plot 5). Compared to Plot 1, Plots 4 and 5 show a similar albeit less prominent tendency, regarding the shift of responsibility from Ag0 to Ag1 , throughout the several approaches to responsibility attribution considered in this paper.

D EXPERIMENTAL RESULTS FOR HP-MIN

In this section, we present the results from Section 5.4 after replacing the HP definition for actual cause with its enhanced version, HP-MIN, which was introduced in Section 5.4.2. As expected, Plots 2a and 6a are identical, since the number of Property 3.3 violations is not affected by whether the contingency minimality condition is satisfied or not. As a result, the differences in Plots 2c and 6c are insignificant. As mentioned in Section 5.4.2, the number of Property 3.6 violations is considerably reduced when replacing HP with HP-MIN. Although, Plot 6e (compared to Plot 2e) shows a similar tendency for the impact on the agents' degrees of responsibility,¹⁵ it can be seen that this impact is still quite large.

E ADDITIONAL COMPARISON RESULTS OF ACTUAL CAUSALITY DEFINITIONS

In this section, we provide some additional empirical insights that we gain by comparing the actual causality definitions from Section 3, in the experimental test-bed of Section 5. Plots 7a and 7b display the number of distinct actual causes and their corresponding size (over all sampled trajectories), respectively. As expected, the BF definition admits a larger number of distinct actual causes, and of greater size than the other two definitions, since it is not equipped with the notion of contingencies. Plot 7c shows the counterfactual improvement admitted by the actual cause-witness pairs of each definition. It can be seen that the HP definition provides pairs that admit greater improvement in general. However, the main reason why this is happening, is because HP allows for non-minimal contingencies (see Section 5.4.2 and Appendix D). Plot 7d validates this intuition.

¹⁵At least for ACCS, ACW and ACW-I which are the only ones being affected.

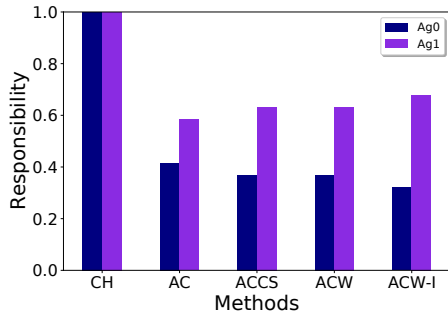


Figure 4: Responsibility Attribution Based on Table 2

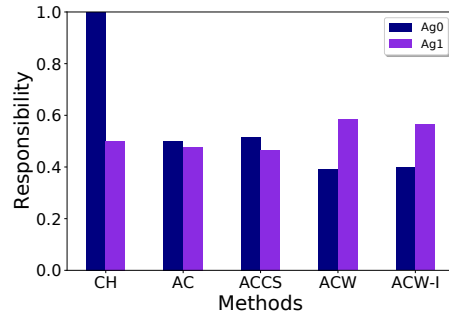
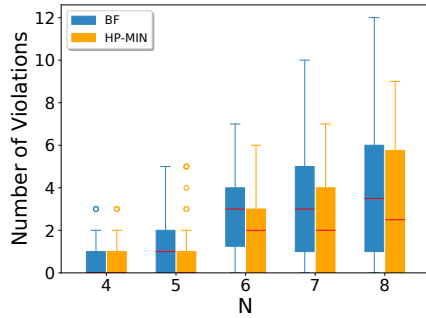
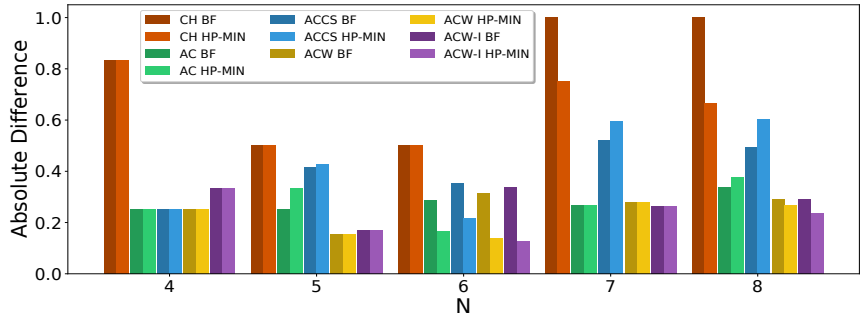


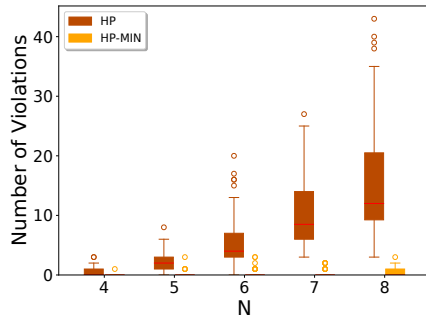
Figure 5: Responsibility Attribution Based on Table 3



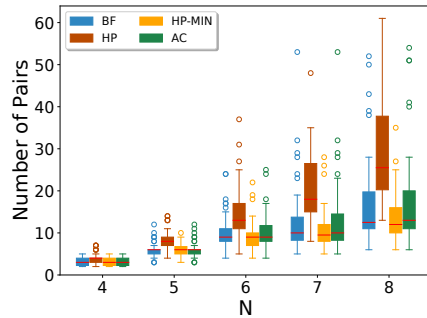
(a) Property 3.3: Number of Violations



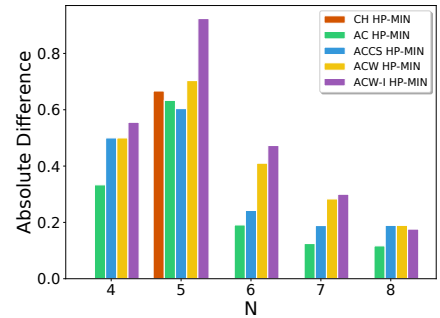
(b) Property 3.3: Impact on Responsibility



(c) Property 3.6: Number of Violations

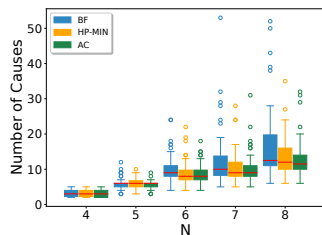


(d) Actual Cause-Contingency Pairs

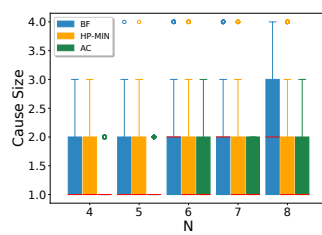


(e) Property 3.6: Impact on Responsibility

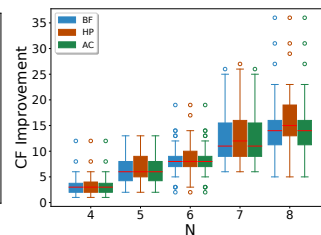
Figure 6: Plots 6a and 6c show the number of violations of Properties 3.3 and 3.6. Plots 6b and 6e show the impact of these violations on the agents' degrees of responsibility. Plot 6d shows the number of distinct actual cause-contingency pairs. Compared to Plots 2a, 2b and 2e, Plots 6a, 6b and 6e have the HP definition replaced by the HP-MIN definition.



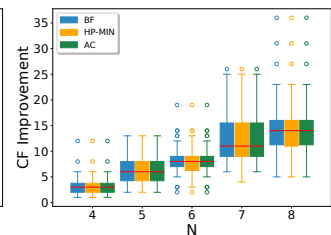
(a) Number of Actual Causes



(b) Size of Actual Causes



(c) CF Improvement



(d) CF Improvement (HP-MIN)

Figure 7: Plots 7a and 7b show the number and size of distinct actual causes. Plots 7c and 7d show the counterfactual improvement admitted by the actual cause-witness pairs of each actual causality definition.