On Blame Attribution for Accountable Multi-Agent Sequential Decision Making

Anonymous Author(s)
Affiliation
Address
email

Abstract

Blame attribution is one of the key aspects of accountable decision making, as it provides means to quantify the responsibility of an agent for a decision making outcome. In this paper, we initiate the study of blame attribution in the context of cooperative multi-agent sequential decision making. As a particular setting of interest, we focus on cooperative decision making formalized by Multi-Agent Markov Decision Processes (MMDP), and we analyze different blame attribution methods derived from or inspired by existing concepts in cooperative game theory. We formalize desirable properties of blame attribution in the setting of interest, and we analyze the relationship between these properties and the studied blame attribution methods. Interestingly, we show that some of the well known blame attribution methods, such as Shapley value, are not performance-incentivizing, while others, such as Banzhaf index, may over-blame agents. To mitigate these value misalignment and fairness issues, we introduce a novel blame attribution method, unique in the set of properties it satisfies, which trade-offs explanatory power (by under-blaming agents) for the aforementioned properties. We further show how to account for uncertainty about agents' decision making policies, and we experimentally: a) validate the qualitative properties of the studied blame attribution methods, and b) analyze their robustness to uncertainty.

... a body of people[1]*, holding themselves accountable to nobody, ought not to be trusted by anybody.*

—Thomas Paine, A philosopher and a political activist.

1 Introduction

With the widespread usage of artificial intelligence (AI) in everyday life [1, 2, 3], accountability has become one of the central problems in the study of AI. Much recent research studied what constitutes accountability in the context of AI and how to design accountable AI systems [4, 5, 6], and recent policies and legislations [7] are increasingly highlighting the importance of accountability, aiming to provide guidelines for developing and deploying accountable AI systems.

Accountability is a relatively broad term, and it typically involves an actor (or multiple actors) justifying their decisions and facing consequences for actions taken [6, 8]. Hence, two critical aspects of accountability are explainability and blame attribution. Recent work proposed various methods for explaining, interpreting, understanding, and certifying algorithmic decision-making and its outcomes [9, 10, 11, 12, 13, 14]. In this paper we study the other critical aspect of accountability – *blame attribution*.

[1] Originally, and by modern standards outdated, Thomas Paine used phrasing with the word *men*.

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

In multi-agent decision making, one of the central roles of blame attribution is assigning blame for undesirable outcomes or, broadly speaking, for the system's inefficiency. Prior work on responsibility and blame in AI [15, 16, 17, 18] has recognized some of the core challenges in attributing blame, including the fact that disentangling agents' contributions to the final outcome is not a trivial task. Such challenges are particularly prominent in sequential settings where past decisions influence the future ones [16].

In this paper, we consider the task of allocating a score to an agent, which represents the degree of its blame, and reflects its contributions to the total inefficiency of the multi-agent system. We focus on cooperative sequential decision making, formalized by multi-agent Markov decision processes (MMDPs) [19], where the outcome of interest is the expected return of the agents' joint policy. Concretely, given an MMDP and the agents' joint policy (true or estimated), we ask: *How to score each agent so that the agents' scores satisfy desirable properties?*

To answer this question, we turn to cooperative game theory and consider blame attribution methods that are derived from or inspired by existing concepts in the cost sharing, data valuation, and coalition formation literature [20, 21, 22, 23, 24, 25, 26, 27], such as core [28], Shapley value [29, 30], or Banzhaf index [31, 32]. Taking this perspective on blame attribution, we initiate the study of blame attribution for accountable multi-agent sequential decision making. More concretely:

- We formalize the set of desirable properties that blame attribution methods should satisfy in cooperative multi-agent sequential decision making. We identify properties that are typically not considered in the cost-sharing literature, yet are important for decision making. In particular, we introduce two novel properties: a) *performance monotonicity*, which states that, having fixed all the other agents, the blame assigned to an agent should not increase if the agent adopts policy that results in a higher expected return (implying that the blame attribution is performance-incentivizing); b) *Blackstone consistency*,[2] which states that an agent should not receive a higher blame just because the agents' policies are not fully known to the blame attribution procedure.
- We characterize the properties of the studied blame attribution methods. We show that some blame assignment methods, such as, Shapley value, are not performance-monotonic (and, hence, performance-incentivizing), while others, such as Banzhaf index, may over-blame agents. Motivated by these results, we introduce a novel blame attribution method that trade-offs explanatory power (by under-blaming agents) for the aforementioned properties.
- We provide algorithms for making the studied blame attribution methods Blackstone-consistent when the agents' policies are estimated. We also characterize the effect of uncertainty on blame attribution methods.
- Using a simulation-based testbed, we experimentally analyze the studied blame attribution methods, their qualitative properties, as well as their robustness to uncertainty. The experiments showcase the importance of the robustness considerations we study and indicate that typically more efficient blame attribution methods (i.e., those that assign more blame in total) are less robust to uncertainty.

1.1 Other Related Work

Apart from the works mentioned in the previous paragraphs, our work relates to different areas of moral philosophy, law, and AI, and here we highlight some of the most relevant references. Research in moral philosophy and law has extensively studied the problem of blame attribution, both in terms of human actors [34, 35, 36], as well as AI actors [37, 38, 39, 40]. We take some of the well known principles in moral philosophy and law in determining properties relevant for blame attribution, e.g., Blackstone consistency is inspired by Blackstone's ratio [33]. In AI, blame attribution has been studied through a more formal lens, utilizing causality [15, 16] and/or game theory [18], and primarily focusing on single-shot decision making and nuances related to defining degrees of responsibility, blame, and blameworthiness. In contrast, we focus on cooperative sequential decision making, and analyze how different blame attribution methods from cooperative game theory fare under different blame attribution properties. Finally, our work is generally related to the *credit assignment* problem [41, 42], and more specifically to the credit assignment problem in multi-agent reinforcement learning [43, 44, 45]. However, our focus is not on supporting the learning processes of agents by reducing computational and statistical challenges of learning, but on evaluating the agents' contributions to the system's inefficiency, ideally in a fair and interpretable manner.

[2] This property is inspired by Blackstone's ratio: "It is better that ten guilty persons escape than that one innocent suffer" [33].

2 Formal Setting

In this section, we describe our formal setting, based on multi-agent Markov decision processes (MMDPs), and we formally model the blame attribution problem in sequential decision making. This section also introduces a set of desirable formal properties of blame attribution methods.

2.1 Preliminaries

We consider a cooperative multi-agent setting, formalized as a class of multi-agent Markov decision processes (MMDPs) $\mathcal{M}$ with n agents $\{1, ..., n\}$. Each MMDP in this class is a tuple $M = (S, \{1, ..., n\}, A, R, P, \gamma, \sigma)$ [19], where: S is the state space; $A = \times_{i=1}^{n} A_i$ is the action space, with A_i being the action space of agent i; R is the reward function $R : S \times A \to \mathbb{R}$ specifying the reward obtained when agents $\{1, ..., n\}$ take a joint action $a = (a^1, ..., a^n)$; P specifies transitions with $P(s, a, s')$ denoting the probability of transitioning to s' from s when agents $\{1, ..., n\}$ take a joint action $a = (a^1, ..., a^n)$; γ is the discount factor; and σ is the initial state distribution. S and A are finite and discrete. A (stationary) joint policy π is a mapping $\pi : S \to \mathcal{D}(A)$, where $\mathcal{D}(A)$ is a probability simplex over A, with $\pi(a|s)$ denoting the probability of taking action a in s. We assume that a joint policy π is factorizable into agents' policies, π_i, i.e., $\pi(a|s) = \pi_1(a_1|s) \cdots \pi_n(a_n|s)$. Therefore, we can define an agent i's policy π_i as a mapping from states to a distribution of agent i's actions, i.e., $\pi_i : S \to \mathcal{D}(A_i)$. We denote the set of all policies by $\Pi = \times_{i=1}^{n} \Pi_i$. We also define a standard performance measure. The return of a joint policy π is defined as $J(\pi) = \mathbb{E}\left[\sum_{t=1}^{\infty} R(s_t, a_t)|s_1 \sim \sigma, \pi\right]$, where the initial state s_1 is sampled from σ, and s_t and a_t are obtained by executing policy π. We abuse our notation by denoting $J(\pi_i, \pi'_{-i}) = J(\pi'_1, ..., \pi_i, ..., \pi'_n)$. Similarly, $J(\pi_S, \pi'_{-S}) = J(\pi'')$ for some $S \subseteq \{1, ..., n\}$, where $\pi''_i = \pi_i$ if $i \in S$ and $\pi''_i = \pi'_i$ if $i \notin S$.

2.2 Blame Attribution

Our goal is to assign blame to agents for failing to jointly achieve optimal performance. Given the agents' behavior policy, denoted by π^b, the inefficiency of the considered multi-agent system can be defined as $\Delta = J(\pi^*) - J(\pi^b)$, where $\pi^* = \arg\max_{\pi} J(\pi)$ is an optimal joint policy. Similarly, we define the marginal inefficiency of a subset of agents $S \subseteq \{1, ..., n\}$ as $\Delta_S = J(\pi_S^{*|b}, \pi^b{}_{-S}) - J(\pi^b)$, as well as the marginal inefficiency of an agent i as $\Delta_i = J(\pi_i^{*|b}, \pi^b{}_{-i}) - J(\pi^b)$, where $\pi_S^{*|b}$ (resp. $\pi_i^{*|b}$) denotes the optimal policy of S (resp. i) assuming all other policies are fixed, i.e., $\pi_S^{*|b} = \arg\max_{\pi_S} J(\pi_S, \pi^b{}_{-S})$ (resp. $\pi_i^{*|b} = \arg\max_{\pi_i} J(\pi_i, \pi^b{}_{-i})$). A blame assignment method is a mapping $\Psi : \mathcal{M} \times \Pi \to R^n_{\geq 0}$, where $\Psi(M, \pi^b)$ distributes blame for inefficiency Δ by assigning score β_i to agent i. The output of Ψ, i.e., the blame vector, is denoted by β.

Uncertainty considerations. Since the agents' behavior policies π^b might not be known to the blame attribution procedure, we also define blame attribution under uncertainty as a mapping $\widehat{\Psi} : \mathcal{M} \times \mathcal{P}(\Pi) \to \mathbb{R}^n_{\geq 0}$ that outputs a blame vector estimate $\widehat{\beta}$. Here, $\mathcal{P}(\Pi)$ represents a set whose elements express the knowledge about π^b. Inspired by the literature on robust MDPs [46, 47, 48], we encode such knowledge with uncertainty sets $\mathcal{P}(\pi^b)$, one associated to each state s, $\mathcal{P}(\pi^b, s)$, defined by the set of probability measures on A. $\mathcal{P}(\pi^b, s)$ identifies the set of plausible stochastic actions that agent i takes in state s. We assume that $\mathcal{P}(\pi^b)$ is consistent with π^b, i.e., $\pi^b(\cdot|s)$ is in $\mathcal{P}(\pi^b, s)$.[3][4]

2.3 Desirable Properties

Our goal is to specify functions Ψ and $\widehat{\Psi}$, such that the assigned scores β_i and $\widehat{\beta}_i$ satisfy desirable properties. In the following text we denote these properties by $\mathcal{R}$. Below, we define properties that are taken from or inspired by the game theory literature [26, 49, 21], but translated to our setting[5]:

- *Validity* ($\mathcal{R}_V$): We say that Ψ is valid if it never distributes more blame than the observed inefficiency Δ. More formally, Ψ satisfies $\mathcal{R}_V$ (resp. ϵ-$\mathcal{R}_V$) if for every M and π^b, $\sum_{i=1}^{n} \beta_i \leq \Delta$ (resp. $\sum_{i=1}^{n} \beta_i \leq \Delta + \epsilon$), where $\beta = \Psi(M, \pi^b)$.

[3] Such $\mathcal{P}(\pi^b)$ can be derived from data containing agents' trajectories and be based on confidence intervals.
[4] $\pi^b(\cdot|s)$ could be in $\mathcal{P}(\pi^b, s)$ w.h.p., provided Blackstone consistency in Section 2.3 is similarly adjusted.
[5] Note that the terminology is slightly different.

- *Efficiency* ($\mathcal{R}_E$): A more strict condition is that the total distributed blame is equal to Δ. That is, Ψ satisfies $\mathcal{R}_E$ (resp. ϵ-$\mathcal{R}_E$) if for every M and π^b, $\sum_{i=1}^n \beta_i = \Delta$ (resp. $|\sum_{i=1}^n \beta_i - \Delta| \leq \epsilon$), where $\beta = \Psi(M, \pi^b)$.
- *Rationality* ($\mathcal{R}_R$): Similar to validity is rationality, which requires that blame distributed to any subset of agents S is not greater than Δ_S. That is, Ψ satisfies $\mathcal{R}_R$ (resp. ϵ-$\mathcal{R}_R$) if for every M, π^b, and $S \subseteq \{1, ..., n\}$, $\sum_{i \in S} \beta_i \leq \Delta_S$ (resp. $\sum_{i \in S} \beta_i \leq \Delta_S + \epsilon$), where $\beta = \Psi(M, \pi^b)$.
- *Symmetry* ($\mathcal{R}_S$): We say that Ψ is symmetric if it treats equal agents equally, i.e., agents that equally contribute to the inefficiency should receive the same blame. More formally, Ψ satisfies $\mathcal{R}_S$ (resp. ϵ-$\mathcal{R}_S$) if for every M and π^b, $\beta_i = \beta_j$ (resp. $|\beta_i - \beta_j| \leq \epsilon$) whenever $\Delta_{S \cup \{i\}} = \Delta_{S \cup \{j\}}$ for all $S \subseteq \{1, ..., n\} \backslash \{i, j\}$, where $\beta = \Psi(M, \pi^b)$.
- *Invariance* ($\mathcal{R}_I$): We say that Ψ is invariant if it assigns zero blame to agents who do not marginally contribute to inefficiency. More formally, Ψ satisfies $\mathcal{R}_I$ (resp. ϵ-$\mathcal{R}_I$) if for every M and π^b, $\beta_i = 0$ (resp. $\beta_i \leq \epsilon$) whenever $\Delta_{S \cup \{i\}} = \Delta_S$ for all S, where $\beta = \Psi(M, \pi^b)$.

Note that $\epsilon > 0$ in the definitions of ϵ-$\mathcal{R}$, and we use these properties used in our characterization result for blame attribution under uncertainty. Additionally, we consider two properties that relate the blame attribution output to the MMDP structure and the agents' behavior policies.

- *Contribution monotonicity* ($\mathcal{R}_{CM}$)[50]: We say that Ψ is contribution-monotonic if the blame it assigns to an agent depends only on its marginal contributions and monotonically so. More formally, Ψ satisfies $\mathcal{R}_{CM}$ (resp. ϵ-$\mathcal{R}_{CM}$) if for every two $(M^1, {\pi^b}^1)$ and $(M^2, {\pi^b}^2)$, ${\beta_i}^1 \geq {\beta_i}^2$ (resp. ${\beta_i}^1 \geq {\beta_i}^2 - \epsilon$) whenever $\Delta^1_{S \cup \{i\}} - \Delta^1_S \geq \Delta^2_{S \cup \{i\}} - \Delta^2_S$ for all S, where $\beta^1 = \Psi(M^1, {\pi^b}^1)$ and $\beta^2 = \Psi(M^2, {\pi^b}^2)$.
- *Performance monotonicity* ($\mathcal{R}_{PerM}$): We say that Ψ is performance-monotonic if it assigns greater blame to agent i for worse performance, assuming the other agents are fixed. More formally, consider any MMDP M, any fixed ${\pi^b}_{-i}$, and any π_i and π'_i such that $J(\pi_i, {\pi^b}_{-i}) \leq J(\pi'_i, {\pi^b}_{-i})$. We say that Ψ satisfies $\mathcal{R}_{PerM}$ (resp. ϵ-$\mathcal{R}_{PerM}$) if $\beta_i \geq \beta'_i$ (resp. $\beta_i \geq \beta'_i - \epsilon$) where $\beta = \Psi(M, (\pi_i, {\pi^b}_{-i}))$ and $\beta' = \Psi(M, (\pi'_i, {\pi^b}_{-i}))$.

The above definitions directly extend to $\widehat{\Psi}$ except that we require them to hold for all $\mathcal{P}(\pi_b)$. Additionally, we identify the following property for $\widehat{\Psi}$:

- *Blackstone consistency* ($\mathcal{R}_{BC}$): We say that $\widehat{\Psi}$ is Blackstone-consistent with Ψ if it never attributes more blame to an agent than Ψ. More formally, $\widehat{\Psi}$ satisfies $\mathcal{R}_{BC}(\Psi)$ if for any M, π^b and $\mathcal{P}(\pi^b)$, $\widehat{\beta}_i \leq \beta_i$, where $\beta = \Psi(M, \pi^b)$ and $\widehat{\beta} = \widehat{\Psi}(M, \mathcal{P}(\pi^b))$.

3 Game-Theoretic Approaches to Blame Attribution

In this section, we study blame attribution methods based on well known game theoretic notions, such as the core [28], Shapley value [29, 30], or Banzhaf index [31, 32]. We also introduce a novel blame attribution method, unique in the set of properties it satisfies. The proofs of our results can be found in the Supplemental Material.

3.1 Max-Efficient Rationality

We start with a relatively simple blame assignment method that puts rationality as a strict condition, and maximizes the efficiency of blame assignment under this constraint. We call this blame assignment method *max-efficient rationality*. More formally, this blame assignment method can be defined via the following linear program:

$$\Psi_{MER}(M, \pi^b) := \max_{\beta} \sum_{i=1}^{n} \beta_i \qquad \text{s.t.} \qquad \sum_{i \in S} \beta_i \leq \Delta_S \quad \forall S \subseteq \{1, ..., n\}, \tag{P1}$$

where Δ_S are precomputed. Max-efficient rationality is inspired by the notion of core, but unlike the core, max-efficient rationality does not require $\mathcal{R}_E$ (efficiency) to hold. It is easy to show that the following properties are satisfied by any optimizer of (P1), Ψ_{MER}.

Proposition 1. *Every solution to the optimization problem (P1), i.e., Ψ_{MER}, satisfies $\mathcal{R}_V$ (validity), $\mathcal{R}_R$ (rationality) and $\mathcal{R}_I$ (invariance).*

In the Supplemental Material we present an algorithm that solves (P1), such that Ψ_{MER} also satisfies $\mathcal{R}_S$ (symmetry). This is also the algorithm that we applied to obtain the results in Section 5. Note that the constraints in (P1) are quite restrictive, leading to blame assignments that typically distribute very little blame in total. The amount of total blame assigned is important for explanatory power. Namely, a trivial blame attribution method that assigns the score of 0 to every agent satisfies all of the properties from the previous section except $\mathcal{R}_E$ (efficiency), but provides no information regarding the agents' contributions to the outcome.

3.2 Marginal Contribution

Another intuitive blame assignment method is what we refer to as *marginal contribution*. This method simply quantifies an agent's potential to increase the performance of the system, assuming that the other agents keep their policies fixed. That is, the blame assigned to agent i is equal to $\beta_i = \Delta_i$. The following properties hold:

Proposition 2. *$\Psi_{MC}(M, \pi^b) = (\Delta_1, ..., \Delta_n)$ satisfies $\mathcal{R}_S$ (symmetry), $\mathcal{R}_I$ (invariance), $\mathcal{R}_{CM}$ (contribution monotonicity) and $\mathcal{R}_{PerM}$ (performance monotonicity).*

Unlike max-efficient rationality, marginal contribution does not satisfy validity, i.e., it can over-blame a group of agents by assigning them total score that exceeds the improvement they can achieve, i.e., Δ. Given that an agent's marginal inefficiency is not always a good indicator of the agent's influence on the system's performance, this method can be highly inefficient (distributing very little) when coordination among agents is required, as we show in in Section 5.

3.3 Shapley Value and Banzhaf Index

In the context of the sequential decision making setting studied in this paper, Shapley value can be defined as $\beta = \Psi_{SV}(M, \pi^b)$ such that:

$$\beta_i = \sum_{S \subseteq \{1,...,n\} \setminus \{i\}} w_S \cdot \left[J(\pi^{*|b}_{S \cup \{i\}}, {\pi^b}_{-S \cup \{i\}}) - J(\pi^{*|b}_S, {\pi^b}_{-S}) \right] \tag{1}$$

where coefficients w_S are set to $w_S = \frac{|S|!(n-|S|-1)!}{n!}$. We restate (and in the Supplemental Material, prove the claim for our setting) a well known uniqueness result for Shapley value:

Theorem 1. *[50] $\Psi_{SV}(M, \pi^b) = (\beta_1, ..., \beta_n)$, where β_i is defined by Eq. (1) and $w_S = \frac{|S|!(n-|S|-1)!}{n!}$, is a unique blame attribution method satisfying $\mathcal{R}_E$ (efficiency), $\mathcal{R}_S$ (symmetry) and $\mathcal{R}_{CM}$ (contribution monotonicity). Additionally, Ψ_{SV} satisfies $\mathcal{R}_V$ (validity) and $\mathcal{R}_I$ (invariance).*

As we show in Section 5, and unlike Ψ_{MC}, Shapley value does not satisfy $\mathcal{R}_R$ (rationality) nor $\mathcal{R}_{PerM}$ (performance monotonicity).

Banzhaf index, denoted by Ψ_{BI}, is similar to Shapley value, and in fact, it has the same functional form but different coefficients ($w_S = \frac{1}{2^{n-1}}$), leading to a slightly different uniqueness result. The Supplemental Material discusses Banzhaf index and its properties in greater detail. Here, we note that Banzhaf index is equivalent to Shapley value for two agents. However, in general, Banzhaf index does not satisfy $\mathcal{R}_E$ (efficiency), but a version of it, called 2-efficiency [51]. As it is the case with Shapley value, Banzhaf index does not necessarily satisfy $\mathcal{R}_{PerM}$ (performance monotonicity) nor $\mathcal{R}_R$ (rationality). Interestingly, $\mathcal{R}_V$ (validity) might also not hold (see Section 5).

3.4 Average Participation

Motivated by the fact that $\mathcal{R}_{PerM}$ (performance monotonicity) is important for incentivizing good performance and $\mathcal{R}_E$ (efficiency) is important for explanatory power, we introduce a novel blame assignment method, which can be seen as a combination of marginal contribution and Shapley value. We first show the following result, which shows that there is an inherent trade-off between $\mathcal{R}_{PerM}$ and $\mathcal{R}_E$, assuming $\mathcal{R}_S$ (symmetry) and $\mathcal{R}_I$ (invariance) hold.

Proposition 3. *No blame attribution method Ψ satisfies $\mathcal{R}_E$ (efficiency), $\mathcal{R}_S$ (symmetry), $\mathcal{R}_I$ (invariance) and $\mathcal{R}_{PerM}$ (performance monotonicity).*

Given this result, we instead consider two new properties $\mathcal{R}_{AE}$ (average-efficiency) and $\mathcal{R}_{cPerM}$ (c-performance monotonicity), which are weaker variants of $\mathcal{R}_E$ (efficiency) and $\mathcal{R}_{PerM}$ (performance

monotonicity), respectively. Importantly, $\mathcal{R}_{AE}$ (average-efficiency) is not satisfied by Ψ_{MC} and $\mathcal{R}_{cPerM}$ by Ψ_{SV}. In addition, we also consider two variants of $\mathcal{R}_{CM}$ (contribution monotonicity), $\mathcal{R}_{cParM}$ (c-participation monotonicity) and $\mathcal{R}_{RcParM}$ (relative c-participation monotonicity). To define the new properties, we introduce a contribution function $c : \mathcal{M} \times \Pi \times \{1, ..., n\}$ that indicates whether an agent marginally contributes to $J(\pi^{*|b}_{S\cup\{i\}}, {\pi^b}_{-S\cup\{i\}})$ for any subset S:[6]

$$c(M, \pi^b, i) = \begin{cases} 0 & \text{if } J(\pi^{*|b}_{S\cup\{i\}}, {\pi^b}_{-S\cup\{i\}}) = J(\pi^{*|b}_{S}, {\pi^b}_{-S}) \quad \forall S \subseteq \{1, ..., N\} \\ 1 & \text{otherwise} \end{cases}.$$

The new properties are then defined as:

- *Average efficiency* ($\mathcal{R}_{AE}$): Ψ satisfies $\mathcal{R}_{AE}$ (resp. ϵ-$\mathcal{R}_{AE}$) if for every M and π^b, $\sum_{i=1}^{n} \beta_i = \sum_{S\subseteq\{1,...,n\}} \frac{1}{2^n-1} \cdot \Delta_S$ (resp. $|\sum_{i=1}^{n} \beta_i - \sum_{S\subseteq\{1,...,n\}} \frac{1}{2^n-1} \cdot \Delta_S| \leq \epsilon$), where $\beta = \Psi(M, \pi^b)$. I.e., compared to efficiency, Δ is replaced by average value of Δ_S.
- *c-Performance monotonicity* ($\mathcal{R}_{cPerM}$): Consider any MMDP M, any fixed ${\pi^b}_{-i}$, and any π_i and π'_i s.t. $J(\pi_i, {\pi^b}_{-i}) \leq J(\pi'_i, {\pi^b}_{-i})$ and $c(M, (\pi_i, {\pi^b}_{-i}), j) = c(M, (\pi'_i, {\pi^b}_{-i}), j)$ for all $j \neq i$. We say that Ψ satisfies $\mathcal{R}_{cPerM}$ (resp. ϵ-$\mathcal{R}_{cPerM}$) if $\beta_i \geq \beta'_i$ (resp. $\beta_i \geq \beta'_i - \epsilon$) where $\beta = \Psi(M, (\pi_i, {\pi^b}_{-i}))$ and $\beta' = \Psi(M, (\pi'_i, {\pi^b}_{-i}))$. I.e., the requirement is that performance monotonicity holds only if the contribution function c does not change for other agents.
- *c-Participation monotonicity* ($\mathcal{R}_{cParM}$): Ψ satisfies $\mathcal{R}_{cParM}$ (resp. ϵ-$\mathcal{R}_{cParM}$) if for every $(M^1, {\pi^b}^1)$ and $(M^2, {\pi^b}^2)$ s.t. $c(M^1, {\pi^b}^1, i) = c(M^2, {\pi^b}^2, i)$ for every i, ${\beta_i}^1 \geq {\beta_i}^2$ (resp. ${\beta_i}^1 \geq {\beta_i}^2 - \epsilon$) whenever $\Delta^1_{S\cup\{i\}} \geq \Delta^2_{S\cup\{i\}}$ for all S, where $\beta^1 = \Psi(M^1, {\pi^b}^1)$ and $\beta^2 = \Psi(M^2, {\pi^b}^2)$. I.e., an agent's blame is larger in an MDP where it has a higher marginal inefficiency.
- *Relative c-participation monotonicity* ($\mathcal{R}_{RcParM}$): Ψ satisfies $\mathcal{R}_{RcParM}$ (resp. ϵ-$\mathcal{R}_{RcParM}$) if for every $(M^1, {\pi^b}^1)$ and $(M^2, {\pi^b}^2)$ s.t. $c(M^1, {\pi^b}^1, i) = c(M^2, {\pi^b}^2, i)$ for every i, ${\beta_j}^1 - {\beta_j}^2 \geq {\beta_k}^1 - {\beta_k}^2$ (resp. ${\beta_j}^1 - {\beta_j}^2 \geq {\beta_k}^1 - {\beta_k}^2 - \epsilon$) whenever $c(M^1, {\pi^b}^1, j) = c(M^1, {\pi^b}^1, k)$ and $\Delta^1_{S\cup\{j\}} - \Delta^2_{S\cup\{j\}} \geq \Delta^1_{S\cup\{k\}} - \Delta^2_{S\cup\{k\}}$ for all S, where $\beta^1 = \Psi(M^1, {\pi^b}^1)$ and $\beta^2 = \Psi(M^2, {\pi^b}^2)$.

Now, we describe the new blame assignment method, which we call, *average participation*. This blame assignment method can be defined as $\beta = \Psi_{AP}(M, \pi^b)$ such that:

$$\beta_i = \sum_{S\subseteq\{1,...,n\}\setminus\{i\}} w \cdot \frac{c(M, \pi^b, i)}{\sum_{j\in S\cup\{i\}} c(M, \pi^b, j)} \cdot \Delta_{S\cup\{i\}}, \tag{2}$$

where $w = \frac{1}{2^n-1}$. Intuitively, this blame assignment method equally distributes blame for inefficiency of $\Delta_{S\cup\{i\}}$ to each of its members that are *pivotal*, i.e., marginally contribute to the inefficiency of some subset of $\{1, ..., n\}$. Agents that are not pivotal, obtain 0 blame. Average participation uniquely satisfies the following properties:

Theorem 2. *$\Psi_{AP}(M, \pi^b) = (\beta_1, ..., \beta_n)$, where β_i is defined by Eq. (2) and $w = \frac{1}{2^n-1}$, is a unique blame attribution method that satisfies $\mathcal{R}_{AE}$ (average-efficiency), $\mathcal{R}_S$ (symmetry), $\mathcal{R}_I$ (invariance), $\mathcal{R}_{cParM}$ (c-participation monotonicity) and $\mathcal{R}_{RcParM}$ (relative c-participation monotonicity). Furthermore, Ψ_{AP} satisfies $\mathcal{R}_{cPerM}$ (c-performance monotonicity) and $\mathcal{R}_V$ (validity).*

Unlike marginal contribution, average participation is valid (never over-blames agents), however it satisfies a weaker version of performance monotonicity. Still, this version is not satisfied by Shapley value. On the other hand, Shapley value is efficient, unlike average participation, which satisfies a weaker requirement—average efficiency. We also showcase these trade-offs in Section 5.

4 Blame Attribution under Uncertainty

In this section, we study blame attribution methods that do not have direct access to π^b. As mentioned in Section 2, we focus on the case where the knowledge about π^b is defined by the uncertainty set $\mathcal{P}(\pi^b)$, and state-wise so that each state is associated with a set of probability measures on A identifying plausible candidates for $\pi^b(\cdot|s)$.[7] We denote $\pi \in \mathcal{P}(\pi^b)$ if π is plausible by $\mathcal{P}(\pi^b)$.

[6] Alternatively, we can say that $c(M, \pi^b, i) = 0$ if and only if agent i's blame in $\Psi_{SV}(M, \pi^b)$ is 0.

[7] Such definition implies a rectangularity of the uncertainty set [46, 48].

4.1 Shapley Value under Uncertainty

In explaining approaches to handling uncertainty, we focus on Shapley value. Arguably, the simplest way to operate under uncertainty is to derive a point estimate of π^b, denoted by $\widehat{\pi}^b$,[8], and apply Ψ_{SV} on this estimate. Albeit being simple, this approach does not satisfy desirable properties, most notably, $\mathcal{R}_V$ (validity) and $\mathcal{R}_{BC}$ (Blackstone consistency).

Validity. Instead of utilizing a point estimate $\widehat{\pi}^b$, we can take $\widehat{\pi}^b$ that lower bounds the total blame assigned to agents. Such $\widehat{\pi}^b$ can be obtained by solving $\widehat{\pi}^b = \arg\max_{\pi \in \mathcal{P}(\pi^b)} J(\pi)$. By applying Eq. (1) with estimate $\widehat{\pi}^b$, we obtain the blame vector that is valid. More formally:

Proposition 4. *$\widehat{\Psi}_{SV,V}(M, \mathcal{P}(\pi^b)) = \Psi_{SV}(M, \widehat{\pi}^b)$ with $\widehat{\pi}^b = \arg\max_{\pi \in \mathcal{P}(\pi^b)} J(\pi)$ satisfies $\mathcal{R}_V$.*

Blackstone consistency. While $\widehat{\Psi}_{SV,V}$ is valid, it might not be Blackstone consistent w.r.t. Ψ_{SV}. In particular, although the total blame is never overestimated, an agent i might receive higher blame than it would receive under Ψ_{SV}. To ensure Blackstone consistency, we can calculate $\widehat{\pi}^{b,i}$ for each agent i separately using the following optimization problem:

$$\widehat{\pi}^{b,i} := \arg\min_{\pi \in \mathcal{P}(\pi^b)} \sum_{S \subseteq \{1,\dots,n\} \setminus \{i\}} w_S \cdot \left[J(\pi^{*|b}_{S \cup \{i\}}, \pi_{-S \cup \{i\}}) - J(\pi^{*|b}_S, \pi_{-S}) \right], \tag{P2}$$

where $w_S = \frac{|S|!(n-|S|-1)!}{n!}$. We obtain the following result:

Proposition 5. *Let $\beta^i = \Psi_{SV}(M, \widehat{\pi}^{b,i})$. Then $\widehat{\Psi}_{SV,BC}(M, \mathcal{P}(\pi^b)) = (\beta^1_1, ..., \beta^n_n)$ satisfies $\mathcal{R}_V$ (validity) and $\mathcal{R}_{BC}(\Psi_{SV})$ (Blackstone consistency w.r.t. $\Psi_{SV}(M, \pi^b)$).*

Note that $\widehat{\Psi}_{SV,BC}$ distributes less total blame than $\widehat{\Psi}_{SV,V}$, since it takes the worst case perspective for each agent separately, while under $\widehat{\Psi}_{SV,V}$ the blame assigned to all agents is computed with the same behavior joint policy. Moreover, the objective function is more complex than in classical robust MDP settings [46, 47], making classical approaches for robust MDPs hard to apply. In practice, we can relax (P2) and optimize a lower bound of the objective; this preserves $\mathcal{R}_{BC}(\Psi_{SV})$ (Blackstone consistency), but at the expense of distributing less blame to the agents. In our experiments, we solve $\min_{\pi \in \mathcal{P}(\pi^b)} J(\pi^{*|b}_{S \cup \{i\}}, \pi_{-S \cup \{i\}})$ and $\max_{\pi \in \mathcal{P}(\pi^b)} J(\pi^{*|b}_S, \pi_{-S})$ for each S separately, and by applying Eq. (1) we obtain the blame assignment.

Other Blame Attribution Methods. Similar approaches also work for other blame assignment mechanisms discussed in Section 3. For example, and focusing on Blackstone consistency, $\widehat{\Psi}_{BI,BC}(M, \mathcal{P}(\pi^b))$ can be obtained in the same way as $\widehat{\Psi}_{SV,BC}(M, \mathcal{P}(\pi^b))$, but with $w_S = \frac{1}{2^{n-1}}$, while $\widehat{\Psi}_{MC,BC}(M, \mathcal{P}(\pi^b))$ can be implemented as $\widehat{\Psi}_{MC,BC}(M, \mathcal{P}(\pi^b)) = (\tilde{\Delta}_1, ..., \tilde{\Delta}_n)$ where $\tilde{\Delta}_i = \min_{\pi \in \mathcal{P}(\pi^b)} J(\pi^{*|b}_i, \pi_{-i}) - J(\pi)$. Implementing Blackstone consistent $\widehat{\Psi}_{MER,BC}(M, \mathcal{P}(\pi^b))$ and $\widehat{\Psi}_{AP,BC}(M, \mathcal{P}(\pi^b))$ is more complex, and we discuss it in the Supplemental Material.

4.2 Characterization Result

Notice that the described Blackstone consistent methods $\widehat{\Psi}(M, \mathcal{P}(\pi^b))$ are not guaranteed to satisfy the properties that their counterparts $\Psi(M, \pi^b)$ satisfy. However, as long as $\widehat{\Psi}(M, \mathcal{P}(\pi^b))$ and $\Psi(M, \pi^b)$ output similar enough blame vectors, properties that hold under $\Psi(M, \pi^b)$ will approximately hold under $\widehat{\Psi}(M, \mathcal{P}(\pi^b))$. More formally, we have the following results.

Theorem 3. *Consider $\widehat{\Psi}$ and Ψ s.t. $\left\| \widehat{\Psi}(M, \mathcal{P}(\pi^b)) - \Psi(M, \pi^b) \right\|_1 \leq \epsilon$ for any M, π^b, and $\mathcal{P}(\pi^b)$. Then if Ψ satisfies a property $\mathcal{R} \in \{\mathcal{R}_V, \mathcal{R}_E, \mathcal{R}_R, \mathcal{R}_S, \mathcal{R}_I, \mathcal{R}_{AE}\}$, $\widehat{\Psi}$ satisfies ϵ-$\mathcal{R}$. Moreover, if Ψ satisfies a property $\mathcal{R} \in \{\mathcal{R}_{CM}, \mathcal{R}_{PerM}, \mathcal{R}_{cPerM}, \mathcal{R}_{cParM}, \mathcal{R}_{RcParM}\}$, $\widehat{\Psi}$ satisfies 2ϵ-$\mathcal{R}$.*

This theorem allows us to quantify the robustness of the blame assignment methods—the closer $\widehat{\Psi}$ is to Ψ, the more robust it is to uncertainty. Interestingly, a trivial blame assignment method that assigns 0 blame to all the agents is robust in this sense. However, as we already mentioned, it is not informative as it does not attribute any blame. This simple example suggests that efficiency (in a broad sense, i.e., how much blame is being distributed) and robustness are sometimes at odds.

[8] For example, this estimate can be derived from data containing the agents' trajectories.

5 Experiments

To demonstrate the efficacy of the studied blame assignment methods, we consider two environments depicted in Fig. 1 and Fig 2. Both environments are adapted from [52] and modified to be multi-agent. The Supplemental Material provides details on the experimental setup and implementation.

Environment 1: This is a Gridworld environment, in which two agents control the same actor but with different priorities. Agent $\mathcal{A}_2$ primarily controls the movement of the actor, while agent $\mathcal{A}_1$ can intervene and override the actions taken by $\mathcal{A}_2$. Cells denoted with S are the initial states, blank cells indicate areas of small negative reward, F cells indicate areas of slightly increased cost and H cells are areas of severe penalty. The cell denoted by G is the terminal state of the environment and has a positive reward. When agent $\mathcal{A}_1$ intervenes in some state, the actor takes the action determined by an optimal single-agent policy, but also pays a cost of intervention C. The behavior policy ${\pi^b}_2$ of agent $\mathcal{A}_2$ is parameterized by variable α, which specifies the probability that $\mathcal{A}_2$ takes an action determined by an optimal single-agent policy, instead of its personal policy. The personal policy of $\mathcal{A}_2$ takes an optimal action with a state dependent probability; otherwise it follows an optimal policy that sets the cost of H cells to the cost of F cells. Agent $\mathcal{A}_1$'s behavior policy ${\pi^b}_1$ optimizes its return and is trained with a model of $\mathcal{A}_2$ specified by the true personal policy of $\mathcal{A}_2$ and variable α' (not necessarily equal to α). In $\mathcal{R}_{cPerM}$ experiments we set $\alpha = 0.4$. In robustness experiments, we only consider uncertainty over the personal policy of $\mathcal{A}_2$ and set $\alpha = 0.2$ and $\alpha' = 0.5$.

S	S	S	S	S	S	S	S
S		F		H			
S			H			F	
S	F				H		F
S			H			F	
S	H	H		F		H	
S	H			H		H	
S			H		F		G

Fig. 1. Gridworld

$\boldsymbol{\mathcal{R}_{PerM}}$: Fig. 3a validates our theoretical results regarding $\mathcal{R}_{PerM}$ (performance monotonicity). More specifically, methods Ψ_{AP} and Ψ_{MC} assign the minimum blame to $\mathcal{A}_1$ when it acts optimally with respect to the true policy of $\mathcal{A}_2$, i.e., when $\alpha' = \alpha$. However, this is not the case for methods Ψ_{SV} and Ψ_{BI}, which implies that these methods are not incentivizing $\mathcal{A}_1$ to act optimally w.r.t. its belief about $\mathcal{A}_2$. Ψ_{MER} and Ψ_{BI} output the same blame assignments as Ψ_{MC} and Ψ_{SV}, respectively.

Environment 2: This is a Graph environment in which 4 agents simultaneously select actions. The graph consists of one starting and one terminal node, as well as 8 intermediate nodes that can be grouped according to their index number; nodes with even index number are located on the upper level of the graph and nodes with odd index number on the lower level. At each time-step every agent chooses to take either action 0 and move to the upper level or action 1 and move to the lower level. We test multiple variants of this environment, each of which defines a different reward function. In all variants, the reward at each time-step is $+1$ if some formation constraint is satisfied and -1 if not. In the first set of experiments (Coordination), we consider 4 different formation constraints: in formation constraint $m \in \{1, ..., 4\}$, at least m agents need to select action 1 for the constraint to be satisfied. Each behavior policy ${\pi^b}_i$ takes action $a_i = 0$ in every node. In the second set of experiments (Robustness), we consider one formation constraint that is satisfied if the agents are arranged equally between the two levels. In states where agents are balanced between the levels, each behavior policy ${\pi^b}_i$ takes the action from the previous time-step with probability p_i; in unbalanced states, the action that leads to the level with the least number of agents is taken with probability p_i.

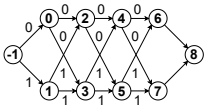

Fig. 2. Graph

Coordination: Fig. 3e shows how much blame in total the blame attribution methods assign for the four different levels of required coordination ($m = 1, ..., 4$). Observe, that when the constraint can be satisfied by every agent ($m = 1$), Ψ_{MC} violates $\mathcal{R}_V$ (validity). For $m = 2$, Ψ_{MER} and Ψ_{MC} assign zero blame to all agents, while Ψ_{BI} violates $\mathcal{R}_V$ (validity). Although always valid, Ψ_{AP} assigns significantly less blame as m increases. Ψ_{SV} is always efficient, and its total blame assignment does not vary with m. Ψ_{SV}, Ψ_{BI} and Ψ_{AP} do not satisfy rationality (as they are greater than Ψ_{MER}).

Robustness: We test the robustness of the blame attribution methods by controlling the amount of uncertainty in the estimates of the agents' behavior policies. To model uncertainty, we consider maximum estimation error ϵ_{max}, and to obtain uncertainty sets $\mathcal{P}(\pi^b)$, we sample (uniformly at random) $\widehat{\pi}^b_i(s)$ such that $\frac{1}{2}\left\|\widehat{\pi}^b_i(s) - {\pi^b}_i(s)\right\|_1 \leq \epsilon_{max}$. Moreover, $\mathcal{P}(\pi^b, s)$ contains all policies π such that $\frac{1}{2}\left\|\widehat{\pi}^b_i(s) - \pi_i(s)\right\|_1 \leq \epsilon_{max}$. In our experiments, we take $\widehat{\pi}^b$ to be the point estimate of π^b.

Comparing estimation approaches: Fig. 3f and 3b show how the approaches for estimating $\widehat{\Psi}_{SV}$ from Section 4 fare under different levels of uncertainty. The point estimate approach typically over-blames an agent and the amount of over-blaming increases with the level of uncertainty. $\widehat{\Psi}_{SV,BC}$ never

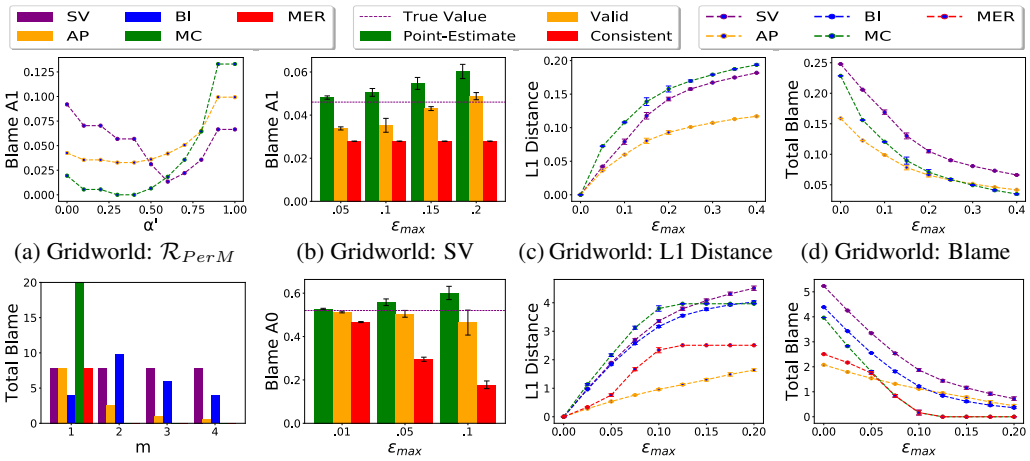

(a) Gridworld: $\mathcal{R}_{PerM}$ (b) Gridworld: SV (c) Gridworld: L1 Distance (d) Gridworld: Blame

(e) Graph: Coordination (f) Graph: SV (g) Graph: L1 Distance (h) Graph: Blame

Fig. 3. Experimental results for the Gridworld and Graph environments. Plot (a) tests methods for $\mathcal{R}_{PerM}$. Plot (e) shows the effect of varying coordination level. Plots (b, c, d, f, g, h) show the effect of varying ϵ_{max} in different Shapley Value approaches (b, f) and blame attribution methods (c, d, g, h).

over-blames any agent, but it becomes less efficient (in distributing blame) as ϵ_{max} increases (Fig. 3f). $\widehat{\Psi}_{SV,V}$ is more efficient than $\widehat{\Psi}_{SV,BC}$, but violates $\mathcal{R}_{BC}$ (Blackstone consistency) (Fig. 3b).

Comparing attribution approaches: Fig. 3g and 3c show for each consistent blame attribution method $\widehat{\Psi}$ from Section 4 the L_1 distance between its output and the output of its counterpart Ψ ("targeted assignment"). Fig. 3h and 3d show the total blame assigned by these methods. $\widehat{\Psi}_{AP,BC}$ consistently outperforms the other methods in terms of the L_1 distance from its "targeted assignment". Compared to $\widehat{\Psi}_{AP,BC}$, $\widehat{\Psi}_{SV,BC}$ is consistently better in terms of efficiency (in distributing blame). Similar, albeit less prominent effects can be seen when comparing $\widehat{\Psi}_{AP,BC}$ and $\widehat{\Psi}_{BI,BC}$. These results indicate a tendency where efficiency (in distributing blame) and robustness are at odds, as we also discuss in Section 4.2. $\widehat{\Psi}_{MER,BC}$ and $\widehat{\Psi}_{MC,BC}$ assign zero total blame even for smaller ϵ_{max}, indicating that they are the least robust to uncertainty.

6 Conclusion

In summary, the focus of our work is to provide an overview of possible computational approaches for attributing blame in multi-agent sequential decision making. We discuss the strength and weaknesses of different methods in order to guide practitioners and policy makers in designing tools that support accountability. We conclude that there is no single best choice for blame attribution methods, since there are inherent trade-offs among properties that one might consider important. Looking forward, we recognize several research directions that could address the limitations of our results; some of which we highlight here. a) In this work we primarily focused on the agents' joint return as the outcome of interest. However, it is often important to pinpoint actual causes that led to more fine grained outcomes. Utilizing a causal perspective would be beneficial in this regard and could link our results to prior work (e.g., [16]). b) We considered model-based approaches to blame assignment. Learning blame attribution directly from data (e.g., with model-free counterfactual RL) might be more practical in settings where an approximate model is hard to obtain. c) More generally, ensuring scalability both in the number of agents and the the richness of environments is one of the most important steps for making this work more widely applicable. We deem approaches from multi-agent RL as suitable candidate solutions for resolving this problem. d) We primarily studied blame assignment properties that are taken from or closely relate to those from the game theory literature. This list could be extended and include more principles from moral philosophy and law. Finally, we would like to draw particular attention to the fact that there is no universal prioritization of properties that applies to all blame attribution problems and hence treating any generic analysis like ours as panacea without further justification, might have a negative impact to the agents that are being blamed. To that end, we would like to emphasize that we see this work, not as a final solution to the blame attribution problem, but as a starting point that shows challenges and trade-offs in distributing blame.

References

[1] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine Bias: There's Software Used Across the Country to Predict Future Criminals. And it's Biased Against Blacks. https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing, 2016.

[2] Amir E Khandani, Adlar J Kim, and Andrew W Lo. Consumer credit-risk models via machine-learning algorithms. *Journal of Banking & Finance*, 34(11):2767–2787, 2010.

[3] Andre Esteva, Brett Kuprel, Roberto A Novoa, Justin Ko, Susan M Swetter, Helen M Blau, and Sebastian Thrun. Dermatologist-level classification of skin cancer with deep neural networks. *nature*, 542(7639):115–118, 2017.

[4] Finale Doshi-Velez, Mason Kortz, Ryan Budish, Chris Bavitz, Sam Gershman, David O'Brien, Kate Scott, Stuart Schieber, James Waldo, David Weinberger, et al. Accountability of ai under the law: The role of explanation. *arXiv preprint arXiv:1711.01134*, 2017.

[5] Joshua A Kroll, Solon Barocas, Edward W Felten, Joel R Reidenberg, David G Robinson, and Harlan Yu. Accountable algorithms. *U. Pa. L. Rev.*, 165:633, 2016.

[6] Maranke Wieringa. What to account for when accounting for algorithms: A systematic literature review on algorithmic accountability. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 1–18, 2020.

[7] European Commission. Ethics Guidelines for Trustworthy Artificial Intelligence. URL: https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai, 2019. [Online; accessed 15-January-2021].

[8] Mark Bovens. Analysing and assessing accountability: A conceptual framework 1. *European law journal*, 13(4):447–468, 2007.

[9] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Model-agnostic interpretability of machine learning. *arXiv preprint arXiv:1606.05386*, 2016.

[10] Anupam Datta, Shayak Sen, and Yair Zick. Algorithmic transparency via quantitative input influence: Theory and experiments with learning systems. In *2016 IEEE symposium on security and privacy (SP)*, pages 598–617. IEEE, 2016.

[11] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 4768–4777, 2017.

[12] Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.

[13] Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*, 2017.

[14] Christoph Dann, Lihong Li, Wei Wei, and Emma Brunskill. Policy certificates: Towards accountable reinforcement learning. In *International Conference on Machine Learning*, pages 1507–1516. PMLR, 2019.

[15] Hana Chockler and Joseph Y Halpern. Responsibility and blame: A structural-model approach. *Journal of Artificial Intelligence Research*, 2004.

[16] Joseph Y Halpern. *Actual causality*. MiT Press, 2016.

[17] Joseph Halpern and Max Kleiman-Weiner. Towards formal definitions of blameworthiness, intention, and moral responsibility. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

[18] Meir Friedenberg and Joseph Y Halpern. Blameworthiness in multi-agent settings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 525–532, 2019.

[19] Craig Boutilier. Planning, learning and coordination in multiagent decision processes. In *Proceedings of the 6th conference on Theoretical aspects of rationality and knowledge*, pages 195–210, 1996.

[20] John Von Neumann and Oskar Morgenstern. *Theory of games and economic behavior (commemorative edition)*. Princeton university press, 2007.

[21] Kamal Jain and Mohammad Mahdian. Cost sharing. *Algorithmic game theory*, 15:385–410, 2007.

[22] Maria-Florina Balcan, Ariel D Procaccia, and Yair Zick. Learning cooperative games. In *Proceedings of the 24th International Conference on Artificial Intelligence*, pages 475–481, 2015.

[23] Eric Balkanski, Umar Syed, and Sergei Vassilvitskii. Statistical cost sharing. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6222–6231, 2017.

[24] Ruoxi Jia, David Dao, Boxin Wang, Frances Ann Hubis, Nick Hynes, Nezihe Merve Gürel, Bo Li, Ce Zhang, Dawn Song, and Costas J Spanos. Towards efficient data valuation based on the shapley value. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1167–1176. PMLR, 2019.

[25] Anish Agarwal, Munther Dahleh, and Tuhin Sarkar. A marketplace for data: An algorithmic solution. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 701–726, 2019.

[26] Yoav Shoham and Kevin Leyton-Brown. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press, 2008.

[27] Georgios Chalkiadakis and Craig Boutilier. Bayesian reinforcement learning for coalition formation under uncertainty. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 3*, pages 1090–1097, 2004.

[28] Donald B Gillies. Solutions to general non-zero-sum games. *Contributions to the Theory of Games*, 4:47–85, 1959.

[29] Lloyd S Shapley. A value for n-person games. *Annals of Mathematics Studies*, (28):307–317, 1953.

[30] Lloyd S Shapley and Martin Shubik. A method for evaluating the distribution of power in a committee system. *The American Political Science Review*, 48(3):787–792, 1954.

[31] John F Banzhaf III. Weighted voting doesn't work: A mathematical analysis. *Rutgers L. Rev.*, 19:317, 1964.

[32] John F Banzhaf III. One man, 3.312 votes: a mathematical analysis of the electoral college. *Vill. L. Rev.*, 13:304, 1968.

[33] William Blackstone and George Sharswood. *Commentaries on the Laws of England. In Four Books*, volume 2. JB Lippincott, 1893.

[34] Thomas M Scanlon. *Moral dimensions*. Harvard University Press, 2009.

[35] David Shoemaker. Attributability, answerability, and accountability: Toward a wider theory of moral responsibility. *Ethics*, 121(3):602–632, 2011.

[36] Ibo Van de Poel, Lambèr Royakkers, and Sjoerd D Zwart. *Moral responsibility and the problem of many hands*. Routledge, 2015.

[37] Mark Coeckelbergh. Artificial intelligence, responsibility attribution, and a relational justification of explainability. *Science and engineering ethics*, 26(4):2051–2068, 2020.

[38] Steve Torrance. Ethics and consciousness in artificial agents. *Ai & Society*, 22(4):495–521, 2008.

[39] Peter M Asaro. Robots and responsibility from a legal perspective. *Proceedings of the IEEE*, 4(14):20–24, 2007.

[40] Gabriel Lima, Nina Grgić-Hlača, and Meeyoung Cha. Human perceptions on moral responsibility of ai: A case study in ai-assisted bail decision-making. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–17, 2021.

[41] Marvin Minsky. Steps toward artificial intelligence. *Proceedings of the IRE*, 49(1):8–30, 1961.

[42] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[43] Kagan Tumer and Adrian Agogino. Distributed agent-based air traffic flow management. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, pages 1–8, 2007.

[44] Jakob Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. Counterfactual multi-agent policy gradients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

[45] Jianhong Wang, Yuan Zhang, Tae-Kyun Kim, and Yunjie Gu. Shapley q-value: A local reward approach to solve global reward games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7285–7292, 2020.

[46] Garud N Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280, 2005.

[47] Arnab Nilim and Laurent El Ghaoui. Robust control of markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.

[48] Aviv Tamar, Shie Mannor, and Huan Xu. Scaling up robust mdps using function approximation. In *International Conference on Machine Learning*, pages 181–189. PMLR, 2014.

[49] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. *Handbook of computational social choice*. Cambridge University Press, 2016.

[50] H Peyton Young. Monotonic solutions of cooperative games. *International Journal of Game Theory*, 14(2):65–72, 1985.

[51] Marcin Malawski. Equal treatment, symmetry and banzhaf value axiomatizations. *International Journal of Game Theory*, 31(1):47–67, 2002.

[52] Cameron Voloshin, Hoang M Le, Nan Jiang, and Yisong Yue. Empirical study of off-policy policy evaluation for reinforcement learning. *arXiv preprint arXiv:1911.06854*, 2019.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] See Sections 2, 3, 4 and 5.
   (b) Did you describe the limitations of your work? [Yes] See Section 6.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section 6.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Section 2 and Section 4.
   (b) Did you include complete proofs of all theoretical results? [Yes] See the Supplemental Material.

3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See the Supplemental Material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 5 and the Supplemental Material.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Section 5 and the Supplemental Material.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See the Supplemental Material.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]