Multi-Agent Influence Diagrams and Commitment

Bachelor Thesis Presented to the
Institute of Computer Science at the
Goethe University Frankfurt

In Partial Fulfillment of the Requirements for the Degree of
Bachelor of Science (B.Sc.)

Author:
Andreas Alexander Haupt

Matriculation Number: 6773522

Supervisor:
Nils Bertschinger

Submitted in August 2019

Erklärung gemäß Bachelor-Ordnung Informatik 2011 §25 Abs. 11
Hiermit erkläre ich, dass ich die Bachelorarbeit selbständig und ohne Benutzung anderer als der angegebenen Quellen und Hilfsmittel verfasst habe. Ebenso bestätige ich, dass diese Arbeit nicht, auch nicht auszugsweise, für eine andere Prüfung oder Studienleistung verwendet wurde.

Frankfurt, 2. September 2019

Andreas Alexander Haupt

Zusammenfassung

Wenn die Interaktion von Umgebung und strategisch handelnden Akteuren komplex ist, ermöglicht ein Multi-Agent Influence Diagram (MAID) [KM03] eine nichtredundante, graphische Repräsentation des strategischen Spiels, die erlaubt, die Struktur des Spiels für schnellere Gleichgewichtsberechnung zu nutzen. MAIDs sind weiterhin gut geeignet, um zu identifizieren, ob Agenten Anreize haben, die Netzwerkstruktur zu verändern bevor sie sich (strategisch) für ihre Aktionen entscheiden [BEW; Eve+19].

Auf einer Menge von MAIDs mit beschränkter Kommunikation definieren wir eine Äquivalenzrelation und geben einen Normalisierungsalgorithmus an. Dieser erlaubt uns, Ergebnisse [BEW] etwas zu verallgemeinern. Weiterhin zeigen wir zwei Möglichkeiten, die optimale Veränderung der Netzwerkstruktur zu berechnen: Eine erste Reduktion auf eine Gleichgewichtsberechnung in einem MAID erlaubt einen endlichen Algorithmus in voller Allgemeinheit. Hingegen zeigen wir für eine Teilklasse von MAIDs eine Darstellung als Lineares Programm, welches in Polynomialzeit lösbar ist.

Abstract

When interaction of the environments and agents is highly structured and complex, MAIDs [KM03] provide a framework for non-redundant, graphical game representations that allows for accelerated equilibrium computation. MAIDs are well-suited to identify incentives for agents to alter the information structure of games before choosing their actions [BEW; Eve+19].

On a set of MAIDs of restricted communication, we define an equivalence relation and present a normalising algorithm. This allows us to slightly generalise results in [BEW]. Furthermore, we show two possibilities for the computation of the optimal network alteration: Our reduction to a MAID equilibrium yields a finite algorithm. For a subclass of MAIDs, we provide a Linear Programming formulation solvable in polynomial time as an alternative.

Contents

List of Contributions

List of Figures

Nomenclature

$\cong_X$	outcome equivalent w.r.t. X
$\neg$	not
$[n]$	$\{1, 2, \dots, n\}$
Θ	Landau notation, bounded above and below asymptotically
D	set of all decision nodes
D^c	set of commitment nodes
d^c	commitment node in liability MAID
d^l_j	liability commitment node in liability MAID
D_i	set of decision nodes of agent i
E	edge set
$p^\frown$	conditional probability mass function of concatenation strategy
p^d	conditional probability mass function of direct mechanism
U	set of all utility nodes
U_i	set of utility nodes of agent i
X	set of nature nodes
$\mathrm{Anc}(A)$	nodes from which A is reachable
$\hat{b}$	auxiliary parent for strategic relevance
$\mathrm{Ch}(A)$	children of A
$\boldsymbol{d}$	marginal distribution of choice rule for $d \in D$
$\text{d-conn}(X, Y, Z)$	d-connection of X and Y given evidence Z
$\Delta(A)$	probability distributions on finite set A
δ_{ij}	Kronecker's delta
$\delta^-(A)$	incoming edges to A
$\delta^+(A)$	out-going edges from A
$\mathrm{Desc}(A)$	nodes reachable from A
δ_t	Dirac distribution
$\mathrm{dom}(\boldsymbol{a})$	domain of random variable
$\text{d-sep}(X, Y \mid Z)$	d-separation of X and Y given evidence Z
$\mathbb{E}$	expected value
$E(A, B)$	directed cut between A and B
$a \dashrightarrow b$	directed a-b-path
$a \to b$	edge $(a, b) \in E$
$a — b$	$(a, b) \in E$ or $(b, a) \in E$
$a \cdots b$	a-b-path in underlying undirected graph
EU_i	expected utility of agent i
$G[V]$	induced subgraph on node set V
$\perp\!\!\!\perp$	stochastic independence
$\mathcal{M}$	generic influence graph
$\mathcal{M}^e$	e-augmentation of influence graph $\mathcal{M}$
$\mathrm{Pa}(A)$	parents of A
$p_{\boldsymbol{X}}$	probability mass function
$\mathcal{R}$	relevance graph of MAID
$\widetilde{u}$	utility function
Unif_A	uniform distribution on finite set A

1 Introduction

Rational agents in a strategic situation try to choose their actions to get the most desirable outcome they can influence the situation to have. In many situations, however, an institution should be able to be freed from these incentives agents face. For example, in an auction, if an auctioneer with negligible interest for an auctioned good receives bids below a previously announced reserve price, she would still want to sell it—if she could not *commit* to a plan of action beforehand.

On the other hand, information from different sources and actions are intertwined in complex ways, necessitating models that capture the structure of interactions. This thesis aims to develop a theory of computing agent actions if they have the power to commit themselves, and information structure is complex. To achieve this, we extend an established model for Bayesian Game theory, MAIDs [KM03], to include decision rules that allow for commitment.

Our extension of the graphical models to include commitment goes in three steps. First, we present possibilities to pre-process representations of strategic interaction to allow for faster computation and equilibrium checks.

Theorem (Main result 1)**.** *There is a quadratic-time algorithm that under the assumption of Conjecture 1 is a canoniser for outcome equivalence on the set of centralised* MAIDs *compatible with a fixed set of chance nodes.*

We use the intuitive concept of outcome equivalence that treats MAIDs as equivalent if any parametrisation of chance/nature node distributions leads to the same expected utility for all agents. This approach is more general than the one typically taken in the literature on complex system design studying the effect of specific network transformation such as the deletion of irrelevant information links [Mil+08], the irrelevant addition of nodes [BEW], and the value of information [How66] (which is zero if a link addition is irrelevant). Our definition allows for the unified treatment of several of the aforementioned questions. For example, if (and only if) the canonical form with respect to outcome equivalence of two MAIDs differing in several added edges is the same, then the addition is irrelevant (but see subsection 3.3).

The second and third contribution then relate to the main problem of computing actions to whose decision rules agents can commit. We first give a reduction to the equilibrium computation in MAIDs for which generic algorithms and approximations exist [MM96].

Theorem (Main result 2)**.** *We can finitely reduce* MAID Equilibrium *to* MAID Commitment Equilibrium*.*

Our transformation only works for a single agent who can commit. In fact, in many modelling situations this is a reasonable assumption: When considering the interaction with potential buyers of one good, an auctioneer might abstract strategic interactions with another auctioneer.

Furthermore, several agents with the power to commit introduce analytical pitfalls. We formulate in the formalism of MAIDs an example from the economics literature illustrating this [Mye82]. Indeed, the example of four agents, two of whom can commit to actions beforehand and choose non-cooperatively, does not permit *any* equilibrium.

Our final contribution shows that given restrictions on agent communication, the agent decisions can be specifically chosen so that they allow for a linear program (LP) formulation.

Theorem (Main result 3)**.** *Finding commitment equilibria in private values* MAIDs *can be reduced to solving a* LP*.*

This result is closely related to formulations of the revelation principle in mechanism design (e.g. [Mye82; Mye86; SW17]) and automatic mechanism design [CS02]. Our approach allows for the *graphical* representation of interactions and to use sparsity to make representations more compact.

We give two applications of our theorems. First, we present additional sufficient properties of a MAID to not increase utility for agents to the ones in [BEW, Proposition 6]. Furthermore, we present a class of MAIDs that permit a polynomial-time solvable LP representation.

The outline of the thesis is as follows: In section 2 we collect notation and give a background on causal graphs and influence diagrams, conditional independence resp. strategic relevance and their connection to d-connectivity and s-reachability which will be used in the rest of the text. In section 3 we give definitions and our normalising algorithm for MAIDs. Sections 4 and 5 present our results on the reduction to MAIDs and LPs, respectively. Finally, we present related literature in section 6 and conclude in section 7.

2 Background

This section gives definitions and reviews known properties of causal graphs and influence diagrams. We start with notations and conventions on graphs, probability distributions and canonical forms in subsection 2.1. Then, we introduce causal graphs in subsection 2.2 and Influence Diagrams (IDs) in subsection 2.3.

We assume prior knowledge of basic probability (conditional probability, marginals, expected value) and graph theory (paths, graph scanning, subgraphs). Prior knowledge in game theory is helpful, but not necessary for this review nor for the rest of the text.

2.1 Preliminaries

Naming Conventions All uppercase letters will be sets (e.g. U, D, A, D_i, U_i, D^c, $\mathrm{Anc}(V)$). The lower-case letters i, j, k, l, n, m mean elements in a non-node set, all other lower-case letters nodes in a graph. All graphs and diagrams will be in script style, e.g. $\mathcal{M}$, $\mathcal{G}$. We will associate nodes in a graph with random variables. If a random variable on domain[1] D is associated to a node v, we write $\mathrm{dom}(\boldsymbol{v}) \coloneqq D$ and denote, depending on context, the random variable or any element in $\mathrm{dom}(\boldsymbol{v})$ by $\boldsymbol{v}$. Similarly, we denote for a set of nodes that jointly have associated random variables V $\mathrm{dom}(\boldsymbol{V}) \coloneqq \bigtimes_{v \in V} \mathrm{dom}(\boldsymbol{v})$. For a fixed enumeration of all nodes of a graph and induced enumeration of a set of nodes $V = \{v_1, v_2, \ldots, v_k\}$ we denote by $\boldsymbol{V}$, depending on context, either the random variable $(\boldsymbol{v_1}, \boldsymbol{v_2}, \ldots, \boldsymbol{v_k})$ or any element of $\mathrm{dom}(\boldsymbol{V}) = \mathrm{dom}(\boldsymbol{v_1}) \times \mathrm{dom}(\boldsymbol{v_2}) \times \cdots \times \mathrm{dom}(\boldsymbol{v_k})$.

As notations, $[n] \coloneqq \{1, 2, \ldots, n\}$. For $U \subseteq V$, we denote by (U, U_C), for $a \in A$, $b \in B$, by $(a, b) \in A \times B$ the concatenation. δ_{ij} is Kronecker's delta.

Probability We assume that all random variables are random variables with respect to one measure space $(\Omega, \mathcal{A}, \mu)$. We denote stochastic independence with respect to this measure space by $\perp\!\!\!\perp$. Let $\Delta(A)$ be the set of probability distributions on set A. For a finite set A, Unif_A is the uniform distribution A, δ_t is the Dirac distribution on A putting probability 1 on $t \in A$. We denote (conditional) probability mass functions by a small p with a subscript for the random variables and the conditioning, e.g. $p_{X|Y}(X|Y)$. If the random variables and the conditioning are clear from the context, we will suppress the subscript.

Graphs All graphs are assumed loopless and simple. For a directed graph $G = (V, E)$, the *corresponding undirected graph* is $(V, \{\{v, w\} | (v, w) \in E \text{ or } (w, v) \in E\}$. For a graph $G = (V, E)$ and $V' \subseteq V$ we use the notation $G[V']$ to denote the induced subgraph of G on V'. We call all nodes in a directed graph with zero in- resp. out-degree *roots* and *leaves*.

We denote by $\delta^-(A)$, $\delta^+(A)$, respectively, the set of all incoming resp. out-going edges. Furthermore, we denote by $\mathrm{Pa}(A)$, $\mathrm{Ch}(A)$, $\mathrm{Anc}(A)$ and $\mathrm{Desc}(A)$, respectively, the sets of all parents, children, ancestors and descendants of node set A. For any of δ^-, Pa, Ch, Anc and Desc, we will omit braces if the argument is a singleton set, e.g. $\delta^-(x) \coloneqq \delta^-(\{x\})$.

For a fixed, implicit edge set E we introduce further notation. We denote $a \to b :\Leftrightarrow (a, b) \in E$. Furthermore, if there is an edge between a and b in the underlying undirected graph, we write $a — b$. Finally, we denote by $a \dashrightarrow b$ a directed path from a to b and by $a \cdots b$ a path in the corresponding undirected graph.

[1] We follow [KM03] in this non-standard nomenclature. In probabilistic terms, the domain should rather be called „range“.

Normal Forms Let $\cong$ be an equivalence relation on set A. Then a function $f\colon A \to A$ such that $f(a) \cong a$ for any $a \in A$ and $f(a) \neq f(b) \to a \ncong b$ for any $a, b \in A$ is called *canoniser*. A function that satisfies only $f(a) \cong a$ is called *normaliser*. The elements of the range of a canoniser are called *canonical forms*, the elements of the range of a normaliser *normal forms*.

2.2 Causal Graphs

Models with uncertainty and interacting entities in high generality are joint probability distributions over a potentially large number of parameters. Indeed, in many settings, even if exact functional relationships are known, measurement errors or variations from outside introduce noise into a model and make point predictions impossible. Moreover, probabilistic modelling allows for a better understanding of the underlying uncertainty.

A (finite) random variable is a (measurable) function from a measure space

$$\boldsymbol{X}\colon (S, \mathcal{A}, \mu) \to \mathrm{dom}(\boldsymbol{X}).$$

We can represent it equivalently by its probability mass function $p_{\boldsymbol{X}}\colon \mathrm{dom}(\boldsymbol{X}) \to [0, 1]$, which for any element in $\mathrm{dom}(\boldsymbol{X})$ specifies the probability weight.

Unfortunately, probability distributions, even when represented via probability mass functions, are highly parametrised. For example, a random variable on binary strings of length 512 bits has a representation size of $2^{2^{512}}$ real numbers. In fact, for a fixed number of states, the number of probability distributions for a growing number of variables grows doubly exponentially. Therefore, to make computational treatment of probability distributions feasible, a lower parametrisation is needed.

Fortunately, the interactions of variables in systems are often extremely sparse. For example, one would assume that there is no relation between the performance in an exam and the food one ate on the first day of lectures. Models that allow for the encoding of sparsity, hence, are essential.

Temporal structure in the sense of a partial order of events is often present in strategic interactions. For example, in a fault prediction scenario or a Q&A system, a question order might underly the system. More generally random trees assume a total temporal ordering of events and are powerful machine learning models [Bis06]. A causal graph is a model that encodes probability distributions of sparse interaction in the presence of a *partial* order of temporal events. In contrast to a random tree, it allows for simultaneous occurrence of events. In more mathematical terms, Bayes nets encode finite, not necessarily exchangeable, relational data.

In addition to causal graphs, we introduce causal models. The difference will become important later, when we formulate purely structural results that only use causal graph structure and quantify over all possible causal models for a given causal graph.

Definition (Causal graph, causal model)**.** *A* causal graph *is a directed acyclic graph (*DAG*)* $\mathcal{G} = (V, E)$*. A* causal model *for* $\mathcal{G}$ *is a set of random variables* $\{\boldsymbol{v}\}_{v \in V}$ *with domains* $\mathrm{dom}(\boldsymbol{v})$ *such that their joint mass function factorises as*

$$p_{\boldsymbol{V}}(\boldsymbol{V}) = \prod_{v \in V} p_{\boldsymbol{v}|\mathbf{Pa}(\boldsymbol{v})}|(\boldsymbol{v}|\mathbf{Pa}(\boldsymbol{v}))$$

We call $\boldsymbol{V}$ *a* parametrisation *of* G*. We will keep the domains of the random variables implicit except where needed. Note that, following our convention, in the subscripts and in the arguments, the same variables have different semantics: In the subscripts, these are random variables, in the argument, they are elements of the domain.*

Causal models allow for much lower parametrised models in contrast to general probability distributions, as they require, for each node, to save only a probability distribution on its domain for each parent instantiation.

The definition of a causal model based on factorisation makes clear that it is a much more economical representation of a probability distribution; the causal structure, i.e. the structure of probabilistic dependence, however, is not clearly identifiable. Characterising probabilistic independencies is crucial for two reasons.

First, knowing the independencies that a model encodes, it is easier to craft a model given expert knowledge. In a medical setting, for example, a doctor might know that some symptoms are likely to be independent of one another, whereas might be interdependent.

Furthermore, the independencies of a model also give insights on model fit. Wrong predicted independencies that contradict expert knowledge in the field of application are a strong signal for misspecification

The following theorem characterises conditional independence of two variables a causal graph

Theorem ([PD96, Theorem 2])**.** *Let* $G = (V, E)$ *be a causal graph. Then for* $X, Y, Z \subseteq V(G)$*,*

$$\boldsymbol{X} \perp\!\!\!\perp \boldsymbol{Y} | \boldsymbol{Z}$$

for any parametrisation $\boldsymbol{V}$ *if and only if*

$$\text{d-sep}\,(X, Y|Z)\,.$$

The (purely graph-theoretical) predicate $\text{d-sep}\,(X, Y|Z)$ *of d-separation is defined as the following: Whenever* $A \cdots B$ *is a path in the corresponding undirected graph, there is an adjacent triple* $u - v - w$ *such that either*

- $u \to w \leftarrow w$*, which we call a* v-structure*, and no descendant of* w *is in* Z*.*
- $u - v - w$ *is not a v-structure and* $v \in Z$*.*

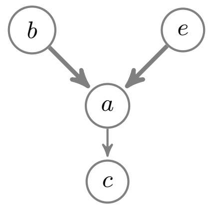

Figure 1: Example of a causal model (cited in [KFB09] as due to Pearl). In this example, either a person is alarmed (a) by an earthquake (e) or a burglary (b). Given that the person is alarmed, she decides whether to make an emergency call (c). Conditional on the fact that there was an emergency call, for some parametrisation, burglary and earthquake are dependent for there is an active path (bold). The v-structure at a descends to evidence c.

We call the negation of d-separation *d-connection*, $\text{d-conn}(X, Y, Z) := \neg\, \text{d-sep}\,(X, Y|Z)$. d-connection is characterised by the existence of a path $X \cdots Y$ that has evidence $z \in Z$ on the path only at v-structures and for which any v-structure descends to evidence. Such paths are called *active*.

For an example of an application of the theorem, see Figure 1. The reader can find a proof of this theorem in [KFB09, Section 3.3.2]. The literature uses the words soundness and completeness for the sufficiency respectively necessity direction in the above theorem. Completeness means that for sets of nodes that are d-connected there are in fact parametrisations that make them stochastically dependent. Soundness is the property that no nodes that are independent under any parametrisation are d-connected.

The purely graph theoretical criterion for independence allows for efficient solution of independence queries. Indeed, d-separation from any fixed node set X is computable by a modified graph scanning in linear time.

As with d-separation, results for IDs, which we introduce next, try to derive properties that hold for any parametrisation given only graphical properties. Soundness results (which say that a criterion is sufficient to show independence) are thus far more widespread than completeness results (which show that the results obtained are best-possible).

2.3 Influence Diagrams

Games are interactions of strategic entities (agents) that, based on information that they are provided, make decisions (actions) to maximise their utility. Utility is assumed to be a number which might depend on all other agents' actions (hence there is a total order on the outcomes for each agent). The goal in (non-cooperative) game theory is to characterise properties of so-called *equilibria* in which an agent does not want to play another action given that all other players stick to their (equilibrium) actions. For an introduction to game theory, see [MWG95, chapter 7].

Technically, a game consists of a finite set of agents $[n]$, and for each agent i an action set D_i and a *utility function*

$$u_i\colon X \times \left(\bigtimes_{i\in[n]} D_i \right) \to \mathbb{R},$$

where X is a set of variables that are not agent actions. The goal for each agent is to maximise his utility function u_i, which depends on his and all other agent's choices. The joint maximisation of utility functions can be seen as an interlinked multi-criteria optimisation problem. Intuitively, action choices are stable if any agent with his choice of action $d_i \in D_i$ maximises u_i given all other agent's action choices $d_j \in D_j$, $j \neq i \in [n]$. With such action choices, no single agent can get a better outcome when he does not cooperate with other agents. Such action choices are called *equilibria*Âăin game theory.

General games are, as probability distributions, highly parametrised. Often, the sets of actions are several independent actions or even actions that depend on other agent's actions, making the set D_i a large product space.

As for probability distributions, however, sparsity in games often occurs. For example, many games consist of sequential moves by several players, often leading to situation in which some players do not take an action anymore: In an incremental auction, as soon as the announced price surpasses their willingness to pay, their actions are irrelevant. Other games have identical but independent interactions by several players and high symmetry properties, such as in voting scenarios.

IDs allow for low-parametrised causal models with a partial time order and have a graphical representation of independence.

Definition ((Multi-Agent) influence graph)**.** *A DAG $\mathcal{M} = (V, E)$ where V partitions as*

$$V = X \cup \bigcup_{i=1}^{n} D_i \cup \bigcup_{i=1}^{n} U_i$$

$U_i, i \in [n]$ only consists of leaf nodes is called influence graph*. We also write $\mathcal{M} = (X, D, U, E)$, keeping the agent sets implicit and using the definitions $D = \bigcup_{i\in[n]} D_i$ and $U = \bigcup_{i\in[n]} U_i$.*

The three different node sets X, A and U encode chance, agent actions and utilities and are typically visualised by different shapes: chance nodes as circles and, filled with different greytones for different players, squares and diamonds for utility nodes. For an example of an influence graph, see Figure 2. The names for the node types vary in the literature: Action nodes are also called *decision* or *choice nodes*, chance nodes also *nature nodes*. Before giving the formal definition of an ID/influence model, parametrisations and equilibria, we describe more thoroughly each of the node types.

First, conditional probability functions parametrise X the set of *chance nodes*—as in causal models. They can encode any physical or even strategic system that the abstract model treats as random.

The second set of nodes, *action nodes* A_i for each agent i, encodes the actions available to an agent. For each action node, an agent makes a (potentially random) decision given (only) realisations of the action node's parents. The selection of a mapping of parent realisation to an outcome is called a *decision rule* and can be identified with a conditional probability distribution (CPD). We highlight that there need not be a total (temporal) order of an agent's nodes nor that agents remember past observations. This allows the formalism to model settings where several independent entities act cooperatively in a strategic environment.

The final set of nodes comes with the largest, despite standard in economics, assumption on agent behavior: The utility nodes are parametrised with deterministic functions of their parent realisations. Given a parametrisation for chance and utility nodes and decision by all agents, one computes an agent's utility by summing over the expected values of all utility nodes. This assumes both that the agents are expected utility maximisers and that their utility function can be *additively* decomposed according to different aspects that influence their utility. Both utility maximisation and additive independence have been axiomatised for general preferences. See [MWG95, Chapter 1] for requirements for expected utility maximisers, [KFB09, Proposition 22.4] for properties for additive separable utility functions.

Definition. *Let $\mathcal{M} = (X, A, U, E)$ be an influence graph. Define:*

Parametrisation *A parametrisation for $\mathcal{M}$ consists of*

- $\operatorname{dom}(\boldsymbol{z})$ *for* $z \in X \cup \operatorname{Pa}(X) \cup \operatorname{Pa}(U)$,
- *For each* $\mathbf{Pa}(\boldsymbol{X}) \in \operatorname{dom}(\mathbf{Pa}(\boldsymbol{X}))$ *a random variable* $\boldsymbol{X}$ *that is a parametrisation for the causal graph* (X, E), *and*
- *Utility functions* $\widetilde{u}\colon \operatorname{dom}(\mathbf{Pa}(\boldsymbol{U})) \to \operatorname{dom}(\boldsymbol{u}) \subset \mathbb{R}$ *for any* $u \in U$. *We can identify* $\widetilde{u}$ *with a conditional probability mass function for random variable* $\boldsymbol{u}$, $p_{\boldsymbol{u}|\mathbf{Pa}(\boldsymbol{u})}(\boldsymbol{u}|\mathbf{Pa}(\boldsymbol{u}))$ *that only puts probability weight on one element of* $\operatorname{dom}(\boldsymbol{u})$.

We call an influence graph together with a parametrisation an influence model *or* influence diagram.

Communication Domain Specification Communication domain specifications *consists of domains* $\operatorname{dom}(\boldsymbol{d})$ *for any* $d \in D \setminus \operatorname{Pa}(X) \cup \operatorname{Pa}(U)$.

Decision Rule *Given a parametrisation and a communication domain specification, a* decision rule *for* $d \in D$ *a domain* $\operatorname{dom}(\boldsymbol{d})$ *together with a mapping*

$$\boldsymbol{d}\colon \operatorname{dom}(\mathbf{Pa}(\boldsymbol{D})) \to \Delta(\operatorname{dom}(\boldsymbol{d})).$$

We can identify a decision rule with a conditional probability mass function

$$p_{\boldsymbol{d}|\mathbf{Pa}(\boldsymbol{d})}(\boldsymbol{d}|\mathrm{Pa}(\boldsymbol{d})).$$

Expected Utility *Given a communication domain specification and decision rules $\boldsymbol{D_i}$, $i \in [n]$ for all agents, we can identify (using the identifications with conditional probability mass functions) $\boldsymbol{X} \cup \boldsymbol{D} \cup \boldsymbol{U}$ with a parametrisation of the causal graph $(X \cup D \cup U, E)$. The dependence of this distribution (and its marginals) we denote by a subscript $_{\boldsymbol{D}}$. Then define the* expected utility *of agent i as*

$$\mathrm{EU}_i[(\boldsymbol{D_i}, \boldsymbol{D} \setminus \boldsymbol{D_i})] \coloneqq \sum_{u \in U_i} \mathbb{E}[\boldsymbol{u}_{(\boldsymbol{D_i}, \boldsymbol{D} \setminus \boldsymbol{D_i})}].$$

We use here the notation $(\boldsymbol{D}_i, \boldsymbol{D} \setminus \boldsymbol{D_i})$ for $\boldsymbol{D} \cup (\boldsymbol{D} \setminus \boldsymbol{D_i})$.

Equilibrium *A communication domain specification $(\mathrm{dom}(\boldsymbol{d}))_{d \in D \setminus \mathrm{Pa}(X) \cup \mathrm{Pa}(U)}$ together with decision rules $\boldsymbol{D}$ for all agents is called an* equilibrium *if for any agent $i \in [n]$*

$$\mathrm{EU}_i[\boldsymbol{D_i}, \boldsymbol{D} \setminus \boldsymbol{D_i}] \geq \mathrm{EU}_i[\boldsymbol{D_i'}, \boldsymbol{D} \setminus \boldsymbol{D_i}] \tag{1}$$

for any decision rules $\boldsymbol{D_i'}$ on the same communication domains. In this case, $\boldsymbol{D_i}$ is a best response *to $\boldsymbol{D} \setminus \boldsymbol{D_i}$. We also say that it* maximises expected utility *given $\boldsymbol{D} \setminus \boldsymbol{D_i}$.*

Attainable Utility Vector *For a given parametrisation the vectors $(\mathrm{EU}_i(\boldsymbol{D}))_{i \in [n]}$ for an equilibrium $\boldsymbol{D}$ (and some communication domain specification) are called* attainable utility vectors.

Let us give an example of the above definitions. We consider the influence graph depicted in Figure 2. The parametrisation we give can be interpreted as a and b using a coordinating signal, $\boldsymbol{s} \in \{0,1\}$ to coordinate their (binary) choice. Each prefers one of the two, but both only get utility if they choose the same value. A full parametrisation for this MAID would consist of

$$\begin{aligned}
\mathrm{dom}(\boldsymbol{s}) &\coloneqq \{0,1\} & \mathrm{dom}(\boldsymbol{u_a}) \coloneqq \mathrm{dom}(\boldsymbol{u_b}) &\coloneqq \{0,1\} \\
\mathrm{dom}(\boldsymbol{a}) &\coloneqq \{0,1\} & \mathrm{dom}(\boldsymbol{b}) &\coloneqq \{0,1\} \\
p_{\boldsymbol{s}}(0) &\coloneqq \frac{1}{2} & p_{\boldsymbol{s}}(1) &\coloneqq \frac{1}{2} \\
\widetilde{u_a}(\boldsymbol{a}, \boldsymbol{b}) &\coloneqq \delta_{\boldsymbol{ab}} & \widetilde{u_b}(\boldsymbol{a}, \boldsymbol{b}) &\coloneqq \delta_{\boldsymbol{ab}}.
\end{aligned}$$

The first line defines domains for the utility nodes, the second for the action nodes must be fixed to ensure that their domains are compatible with the utility functions' signature. In this example, one does not have to specify communication domains, as $X \cup U \cup \mathrm{Pa}(X) \cup \mathrm{Pa}(X) \supseteq D$. An example of an equilibrium is given by the decision

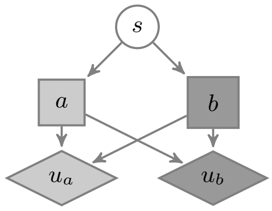

Figure 2: An example of an influence diagram, adapted from [Mil+08, p.2].

rules that copy the value of $\boldsymbol{s}$:

$$
\begin{array}{llll}
p_{\boldsymbol{a}|\boldsymbol{s}}(1|0) = 0 & p_{\boldsymbol{a}|\boldsymbol{s}}(1|1) = 1 & p_{\boldsymbol{a}|\boldsymbol{s}}(0|1) = 0 & p_{\boldsymbol{a}|\boldsymbol{s}}(0|0) = 1 \\
p_{\boldsymbol{b}|\boldsymbol{s}}(1|0) = 0 & p_{\boldsymbol{b}|\boldsymbol{s}}(1|1) = 1 & p_{\boldsymbol{b}|\boldsymbol{s}}(0|1) = 0 & p_{\boldsymbol{b}|\boldsymbol{s}}(0|1) = 1
\end{array}
$$

This equilibrium gives each agent (we do the calculation for u_a, which we assume belongs to agent 1, as both agents are symmetric):

$$
\begin{aligned}
\mathrm{EU}_i[(\boldsymbol{a}, \boldsymbol{b})] &= \sum_{\boldsymbol{s},\boldsymbol{a},\boldsymbol{b},\boldsymbol{u_a} \in \{0,1\}} p_{(s,a,b,u_a)} u_a \\
&= p_{\boldsymbol{s}}(1) p_{\boldsymbol{a}|\boldsymbol{s}}(1|1) p_{\boldsymbol{b}|\boldsymbol{s}}(1|1) p_{\boldsymbol{u_a}|\boldsymbol{a},\boldsymbol{b}}(1|1,1) \cdot 1 \\
&\quad + p_{\boldsymbol{s}}(0) p_{\boldsymbol{a}|\boldsymbol{s}}(0|0) p_{\boldsymbol{b}|\boldsymbol{s}}(0|0) p_{\boldsymbol{u_a}|\boldsymbol{a},\boldsymbol{b}}(1|0,0) \cdot 1 \\
&= \frac{1}{2} + \frac{1}{2} = 1
\end{aligned}
$$

It is not too hard to show that the described decision rules are an equilibrium and that the set of attainable utility vectors for this parametrisation is $\{(\frac{1}{2}, \frac{1}{2}), (1,1)\}$.

The example already shows that in MAIDs complicated structures of coordination might emerge due to shared observations of chance variables. The question of finding an equilibrium for a parametrised influence graph is a formidable one.

MAID Commitment Equilibrium

Input A parametrised MAID $\mathcal{M}$
represented by lists of real numbers $(\widetilde{u}(\mathbf{Pa}(\boldsymbol{u})))_{\mathrm{Pa}(\boldsymbol{u}) \in \mathrm{dom}(\mathbf{Pa}(\boldsymbol{u}))}$, $u \in U$, and $(p(\boldsymbol{x}|\mathbf{Pa}(\boldsymbol{x})))_{\substack{\mathbf{Pa}(\boldsymbol{x}) \in \mathrm{dom}(\mathbf{Pa}(\boldsymbol{x})) \\ \boldsymbol{x} \in \mathrm{dom}(\boldsymbol{x})}}$, $x \in X$
(domains for $\mathbf{U} \cup \mathbf{X} \cup \mathrm{Pa}(\mathbf{U}) \cup \mathrm{Pa}(\mathbf{X})$ are implicitly represented in the lists)

Output Decision rules $\boldsymbol{D}$
represented by lists of real numbers $(p(\boldsymbol{d}|\mathbf{Pa}(\boldsymbol{d})))_{\substack{\mathbf{Pa}(\boldsymbol{d}) \in \mathrm{dom}(\mathbf{Pa}(\boldsymbol{d})) \\ \boldsymbol{d} \in \mathrm{dom}(\boldsymbol{d})}}$, $d \in D$
(domains for $\boldsymbol{D} \setminus (\boldsymbol{U} \cup \boldsymbol{X} \cup \mathbf{Pa}(\boldsymbol{U}) \cup \mathbf{Pa}(\boldsymbol{X}))$ implicitly represented in the lists)

Goal	Choose the decision rules such that they form an equilibrium.

If an upper bound on the cardinality of dom($\boldsymbol{D} \setminus (\boldsymbol{U} \cup \boldsymbol{X} \cup \mathbf{Pa}(\boldsymbol{U}) \cup \mathbf{Pa}(\boldsymbol{X}))$) can be given, finite-time exact and polynomial time approximation schemes can be applied to this problem (see e.g. [MCZ12]. In higher generality, it is not clear how to choose these domains. We will come back to this subject in section 5.

Although no worst-case bound can be obtained, the computation of equilibria can be sped up considerably by making the network sparser if some dependencies are irrelevant or by dividing the problem into manageable sub-parts.

The first approach is best described by criteria for edge deletion.

Proposition ([Mil+08]). *Let $\mathcal{M} = (X, A, U, E)$ be an influence graph and $\{v, w\} \in E$, $w \in D$. If*

$$\text{d-sep}\left(v, \text{Desc}(v) \cap U_i | \text{Pa}(w) \cup \{w\}\right),$$

then for any parametrisation, any attainable expected utility vector in $\mathcal{M} = (X, A, U, E \setminus \{e\})$ is attainable in $\mathcal{M}$.

This result is in contrast to the d-separation criterion only a sufficient criterion. It shows that some of the information is ignorable but not that deleting an edge that does not satisfy the criterion *does* change the set of attainable utility vectors. Section 3 will work towards an if-and-only-if that was begun for single-agent single-action IDs in [Eve+19] by providing further normalisation for a fundamental study of canonical forms of MAIDs.

Another speedup of equilibrium computation can be obtained by divide-and-conquer and using a notion of *strategic relevance.* Using this structure, we can decompose a game into smaller, independent sub-games.

Definition. *Let $\mathcal{M}$ be an influence graph and $a, b \in D$. Then a* strategically relies *on b if the following holds:*

There are two sets of decision rules $(\boldsymbol{d})_{d \in D}$ and $(\boldsymbol{d}')_{d \in D}$ such that

$$\boldsymbol{d} = \boldsymbol{d}', d \neq b$$

and the decision rule $\boldsymbol{a}$ associated to a that maximises expected utility given $\boldsymbol{d} \setminus \{\boldsymbol{b}\}$. Given this, there is no $\boldsymbol{a}'$ such that

$$p(\boldsymbol{a}|\mathbf{Pa}(\boldsymbol{a})) = p(\boldsymbol{a}'|\mathbf{Pa}(\boldsymbol{a}))$$

for any $\mathbf{Pa}(\boldsymbol{a})$ such that $p(\mathbf{Pa}(\boldsymbol{a})) > 0$ that is utility-maximising given $(\boldsymbol{d}')_{d \in D \setminus \{a\}}$.

Less technically, there are strategies for all other decision nodes such that a change only at $\boldsymbol{b}$ will imply that a best response $\boldsymbol{a}$ loses the status of being a best response.

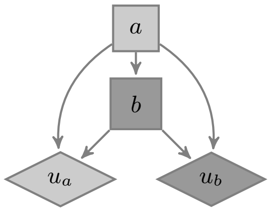

Figure 3: Spurious strategic relevance for zero-probability parent instantiations.

The last part in the definition is to rule out spurious dependencies, compare Figure 3. We take the parametrisation $\mathrm{dom}(\boldsymbol{u_a}) = \mathrm{dom}(\boldsymbol{u_b}) = \mathrm{dom}(\boldsymbol{a}) = \mathrm{dom}(\boldsymbol{b}) = \{0, 1\}$, $\widetilde{u_a}(\boldsymbol{a}, \boldsymbol{b}) = U_{\boldsymbol{a}}(\boldsymbol{a}, \boldsymbol{b}) = \delta_{\boldsymbol{ab}}$.

Independent of the decision rule $\boldsymbol{a}$, a best response is $p(\boldsymbol{b}|\boldsymbol{a}) = \delta_{\boldsymbol{ab}} = \delta_{\boldsymbol{ab}}$ i.e. that $\boldsymbol{b}$ is chosen to match $\boldsymbol{a}$. Given $p(\boldsymbol{a}) = \delta_{\boldsymbol{a}1}$, i.e. $\boldsymbol{a}$ is 1 with certainty, $p(\boldsymbol{b}|\boldsymbol{a}) = \delta_{\boldsymbol{b}0}$ is also a best response but no best response to any other choice of decision rule $\boldsymbol{a}$, which is not too hard to see. This does not imply strategic reliance of $\boldsymbol{b}$ on $\boldsymbol{a}$ as

$$\delta_{\boldsymbol{b}0} \neq \delta_{\boldsymbol{ab}}$$

holds only for $\boldsymbol{a} = 1$, but $p(\boldsymbol{a}) = 0$ if $p(\boldsymbol{a}) = \delta_{\boldsymbol{a}1}$. Similar to d-connectedness and the characterisation of independence for causal graphs, [KM03] showed an analogous necessary and sufficient graphical criterion.

Theorem. *Let $\mathcal{M} = (X, D, U, E)$ be an influence graph. There is a parametrisation for $\mathcal{M}$ such that $\boldsymbol{a} \in D_i$ strategically relies on $\boldsymbol{b} \in D$ if b is s-reachable from a, i.e.: If one adds an auxiliary parent $\hat{b}$ to b, then by treating the directed graph M as causal graph*

$$\text{d-conn}(\hat{b}, U_i, a \cup \mathrm{Pa}(a))).$$

In Figure 3, $\hat{b} \to b \to u_a$ is an active path certifying that a strategically relies on b; for the other direction, the only path is blocked by b. For more examples for s-connectivity, we refer the reader to [KM03, p.206 bottom]. s-reachability defines a relation that we can interpret as a directed graph, the so-called *relevance graph* $\mathcal{R}(\mathcal{M})$. Intuitively, decision rules in strongly connected components can be chosen independently. This intuition is confirmed: One can decompose equilibrium computation in a divide-and-conquer manner to speed up equilibrium computation.

Theorem ([KM03, Theorem 6.2])**.** *Let $\boldsymbol{M}$ be a* MAID *where each agent has perfect recall*[2]*. Then an algorithm that computes equilibria for strongly connected components of the relevance graph separately and in reverse topological order in the component graph of $\mathcal{R}(\mathcal{M})$ outputs an equilibrium for the whole* MAID *when all not-yet-computed decision rules have arbitrary full-support mass functions.*

[2] An agent i is said to have *perfect recall* if there is a total order $d_1, d_2, \ldots, d_k$ on D_i such that $\mathrm{Pa}(d_i) \subseteq \mathrm{Pa}(d_{i+1})$, i.e. an agent recalls all previous actions according to some temporal order.

This theorem makes the size and number of strongly connected components of the relevance graph an important factor for speed-ups of equilibrium computations. We will return to this in section 4.

After this introduction to causal graphs and influence diagrams, we start in the next section by studying *equivalence* of MAIDs. After introducing our notion of equivalence, we slightly generalise a results from [BEW].

3 Normalisation of Multi-Agent Influence Diagrams

This section introduces a unified framework for the comparison of MAIDs. Indeed, deviating from prior literature, our approach does not analyse specific operations that preserve attainable utility vectors but aim for canonical representations given an equivalence relation. The outline of the section is as follows: We start with a definition of equivalence of MAIDs. Then, we present a normalisation algorithm on centralised MAIDs. We conclude this section with a slight generalisation of results in [BEW] and a discussion of the limits of the approach.

3.1 Outcome Equivalence

To define an equivalence of MAIDs on firm ground, we first need the term of a centralised, compatible MAID.

Definition (Centralised, compatible MAIDs)**.** *Let $\mathcal{M} = (X, D, U, E)$ be an influence graph. $\mathcal{M}$ is called* centralised *if*

$$E(D_i, D_j) = \emptyset$$

for any $i, j \in [n]$. It is called compatible *with a* DAG *$\mathcal{G}$ together with in-degrees $(d_v)_{v \in G}$ if $\mathcal{G}$ is a subgraph of $(X \cup D \cup U, E)$, $|\text{Pa}_{\mathcal{M}}(v)| = d_v$ for any $v \in V(G)$ and $X = V(\mathcal{G})$.*

Denote $\mathcal{A}_{(\mathcal{G},(d_v)_{v \in V(G)})}$ the set centralised MAIDs *compatible with $(\mathcal{G}, (d_v)_{v \in V(G)})$ such that $X = V(G)$ and there is exactly one player* 0 *that does not possess any utility nodes.*

The notion of compatibility is needed to ensure that quantifying over all parametrisations is well-defined, as these also define domains for parents of chance and utility nodes. The requirement of centrality and the existence of an indifferent agent will become clearer in the algorithm.

Definition (Outcome equivalence)**.** *We call $\mathcal{M}_1, \mathcal{M}_2 \in \mathcal{A}_X$* outcome equivalent *if for any parametrisation of X the set of attainable expected utility vectors is the same in M_1 and M_2. In this case we write $M_1 \cong_X M_2$.*

Evidently, $\cong_X$ defines an equivalence relation and hence canoniser exist. We conjecture that the algorithm given next is a canoniser for this equivalence relation. We highlight that this algorithm is purely graph-theoretical and does not use any information on parametrisation.

3.2 Normalisation Algorithm

Algorithm 1: Normalisation for $\mathcal{M}$-centralised MAIDs

---

Data: MAID $\mathcal{M}$

Result: Equivalent MAID $\mathcal{M}$

Let $v_1, v_2, \ldots, v_k$ be a topological order of D

1 **for** $i = 1$ **to** k **do** //Information Down

  if $\mathrm{Ch}(v_i) \subseteq D$ *and* $\mathrm{Pa}(v_i) \cap X \neq 0$ **then**

    $E \longleftarrow E \setminus ((\mathrm{Pa}(v_i) \cap X) \times \{v_i\}) \cup ((\mathrm{Pa}(v_i) \cap X) \times \mathrm{Ch}(v_i))$

2 **for** $i = k$ **downto** 1 **do** //Coordination Up

  if $\mathrm{Ch}(v_i) \subseteq D$ *and* $\mathrm{Pa}(v_i) \neq \emptyset$ **then**

    $E \longleftarrow E \setminus (\delta^-(v_i) \cup \delta^+(v_i)) \cup (\mathrm{Pa}(v_i) \times \mathrm{Ch}(v_i))$

    $D \longleftarrow D \setminus \{v_i\}$

3 Find a fixed point of

  $E \longleftarrow E \setminus \{(v, w) \in \delta^-(D) | \text{ d-sep}\,(v, U \cap \mathrm{Desc}(w) | \{w\} \cup \mathrm{Pa}(w))\}$

4 Let $C_1, C_2, \ldots, C_l$ be connected components in

  $(D, \bigcup_{u \in U} \binom{\mathrm{Anc}(u)}{2} \cap \bigcup_{d \in D : \mathrm{Pa}(d) = \emptyset, \mathrm{Ch}(d) \subseteq D} \binom{\mathrm{Ch}(d)}{2}))$

  $E \longleftarrow E \setminus (\bigcup_{d \in D : \mathrm{Pa}(d) = \emptyset} \delta^+(d))$

  $V \longleftarrow V \setminus \{d \in D | \mathrm{Pa}(d) = \emptyset\} \cap \{c_1, c_2, \ldots, c_l\}$

  $E \longleftarrow E \cup (\bigcup_{i \in [l]} \{c_i\} \times C_i)$

5 **for** $i \in [n]$ **do**

  for $v \in D_i$ **do**

    if d-sep(v, U_i) **then**

      $D_i \longleftarrow D_i \setminus \{v\}$

      $D_0 \longleftarrow D_0 \cup \{v\}$

return $\mathcal{M}$

Theorem 1 (MAID normaliser)**.** *Algorithm 1 is a quadratic-time algorithm that is an equivalence transformation for $\cong_X$. Under the assumption that Conjecture 1 holds, it is a canoniser for $(\mathcal{A}_X, \cong_X)$.*

We comment on the lines of the algorithm. The first two lines use that the communication domains are part of an equilibrium specification and can hence be chosen. In this way, much of the communication of action nodes of the same agent can be simplified. Line 1 moves information edges downstream to lower nodes as long as these do not have chance or utility nodes as children. Then, in line 2, all nodes all whose parents and children are decision nodes of the same agent are deleted and their parents are connected to their children.

This graph already has the property that the decision nodes of agents are partitioned into two classes of nodes: Those that have zero in-degree (which we call *coordination devices*) and those that have at least one child in X (which we call *input nodes*).

Line 3 deletes irrelevant edges in the graph. Line 4 determines whether, for any two children, the MAID putting them as children of different coordination devices gives an equivalent graph and if so, splits the coordination device. Finally, line 5 assigns all nodes that are irrelevant to the agent controlling them to agent 0, which is assumed to have constant utility independent of chance and decision rule instantiations.

The following two results, the second of which we are unfortunately unable to prove, would imply that algorithm 1 is a canoniser for outcome equivalence.

Lemma 1 (Equivalence)**.** *Each of lines 1 to 5 in algorithm 1 is an equivalence transformation of M w.r.t. $\cong_X$.*

Conjecture 1 (Uniqueness)**.** *For all outputs of algorithm 1, graph equality and outcome equivalence coincide.*

Establishing Conjecture 1 is challenging as completeness results for MAIDs concerning the irrelevance of edges are not known. Such a result, however, is promised for a follow-up paper of [Eve+19].

Proof of Theorem 1. Correctness follows from Lemma 1 and Conjecture 1. For the runtime note that all lines except for line 3 and line 4 can easily be implemented in quadratic time. For the first, use the algorithm in [KM03] with the modification that each utility nodes replaced by n clones of utility nodes for each agent except agent 0. For the second, observe that the graph construction takes quadratic time (and has size linear in the original graph). Indeed, the first part of the node set can be constructed by at most $|V(M)|$ passes of graph scanning, the latter part is of linear complexity. Finding the connected components is linear. Therefore, this line has a quadratic complexity as well. □

Proof of Lemma 1. For the first two lines of the algorithm, for any iteration in the respective for loop, we will give an alternative definition of communication domains and decision rules that induce the same distributions on decision nodes. For this, we will denote by a prime that a random variable has been transformed, e.g. $\boldsymbol{d}$ is transformed to $\boldsymbol{d'}$.

Define the communication domain of the transformed variable as

$$\operatorname{dom}(\boldsymbol{d'}) := \operatorname{dom}(\boldsymbol{d})^{\operatorname{dom}(\mathbf{Pa}(\boldsymbol{d}) \cap \boldsymbol{X})}$$

and let

$$p(\boldsymbol{d'}|\mathbf{Pa}(\boldsymbol{d}) \cap \boldsymbol{D}) = (p(\boldsymbol{d}|\mathbf{Pa}(\boldsymbol{d}) \cap \boldsymbol{X}, \mathbf{Pa}(\boldsymbol{d}) \cap \boldsymbol{D}))_{\mathbf{Pa}(\boldsymbol{d})\cap\boldsymbol{X}\in\mathrm{dom}(\mathbf{Pa}(\boldsymbol{d})\cap\boldsymbol{X})}.$$

Hence d' outputs a random variable for each potential realisation $\mathbf{Pa}(\boldsymbol{d}) \cap \boldsymbol{X}$ given the distribution of $\boldsymbol{d}$. From this, the child strategies (which are burdensome to formulate) take the values of $\mathbf{Pa}(\boldsymbol{d}) \cap \boldsymbol{X}$ together with the random variable of d' to „look up" the right random variable. This transformation induces the same distribution on the child nodes of d. On the other hand, given p', one can define

$$\mathrm{dom}(\boldsymbol{d}) \coloneqq \mathrm{dom}(\boldsymbol{d'}) \times \mathrm{dom}(\mathbf{Pa}(\boldsymbol{d'}))$$

having a decision rule that concatenates the random variables $\boldsymbol{p'}$ and $\mathbf{Pa}(\boldsymbol{p'})$. Then, the children of p can extract the two separate parts from the concatenated message. This also gives the same joint distribution on the children and parents of d, which is sufficient to establish equivalence, as $\boldsymbol{d}$ is irrelevant for utility given its parents and children.

In line 2, we replace a node d that has only action nodes in its neighbourhood by edges from any parent to any child. Note that by line 1, all of the parents of d have no child in $X \cup U$, we can hence define communication domains.

Choose one parent v of d. Define

$$\mathrm{dom}(\boldsymbol{v'}) \coloneqq \mathrm{dom}(\boldsymbol{v}) \times \mathrm{dom}(\boldsymbol{d})^{\mathrm{dom}(\mathbf{Pa}(\boldsymbol{d}))}$$

and let v simulate for any parent realisation $\mathrm{dom}(\mathbf{Pa}(\boldsymbol{d}))$ in the untransformed random variable the outcome. The children of d then, similarly to as we showed in line 1, „look up" the random variable that they use. This construction implies the same joint distribution of child nodes of d. This guarantees equivalence as after the preprocessing in line 1, only the child nodes are utility relevant. On the other hand, define

$$\mathrm{dom}(\boldsymbol{d}) \coloneqq \mathrm{dom}(\mathrm{Pa}(\boldsymbol{d'})),$$

the concatenation of all parent realisations. Again, d's children can use only the relevant part in the concatenated message. As this also implies the same joint distribution on the children of d, we proved that also line 2 is an equivalence transformation.

Concerning line 3, observe that by means of s-connectivity, the random variable $\boldsymbol{U}$ is independent of v given the parents of w and w itself. Hence

$$E[\boldsymbol{U}|\boldsymbol{w}, \mathbf{Pa}(\boldsymbol{w}), \boldsymbol{v}] = E[\boldsymbol{U}|\boldsymbol{w}, \mathbf{Pa}(\boldsymbol{w}), \boldsymbol{v'}].$$

Therefore, the decision rule at w cannot depend on v and the edge can be deleted without changing expected utility in any setting.

In line 4, some of the coordination devices are split up. A split happens if for a partition of child nodes of a device, no nodes in different partition sets have a common descendant utility node. To show that such a split cannot change expected utility, we consider an auxiliary construction: For a coordination device d and a partition of its child set $C_1 \cup C_2 \subseteq \mathrm{Ch}(d)$, replace d by two connected nodes $d_1 \to d_2$ that have child sets $\mathrm{Ch}(d_1) = C_1$ and $\mathrm{Ch}(d_2) = C_2$ (this construction is an inverse operation as in line 2 and hence an equivalence transformation). It is not too hard to see that if the edge $d_1 \to d_2$ can be deleted as an equivalence transformation, the split of d is also an equivalence transformation. By assumption, there are no directed paths $C_1 \dashrightarrow U \dashleftarrow C_2$. Hence as d_2 has no parents but d_1,

$$\text{d-sep}\left(d_1, U | \{d_2\} \cup \mathrm{Pa}(d_2)\right).$$

Therefore,

$$\mathbb{E}[\boldsymbol{U} | \{\boldsymbol{d_2}\} \cup \mathbf{Pa}(\boldsymbol{d_2}), \boldsymbol{d_1}] = \mathbb{E}[\boldsymbol{U} | \{\boldsymbol{d_2}\} \cup \mathbf{Pa}(\boldsymbol{d_2}), \boldsymbol{d_2}],$$

and $\boldsymbol{d_1}$ is irrelevant to the choice of $\boldsymbol{d_2}$ concerning utility of any agent. Therefore, also line 4 is an equivalence transformation.

Finally, consider line 5. Let $v \in D$. Any decision rule that was a best response when v belonged to D_i, $i \in [n]$ is also a best response when $v \in D_0$. Indeed, agent 0 is indifferent between all outcomes. Conversely, we need to show that agent i is indifferent between any choice of $\boldsymbol{v}$. The condition d-sep(v, U_i) implies that $\boldsymbol{v} \perp\!\!\!\perp \boldsymbol{U_i}$, hence

$$E[\boldsymbol{U_i} | \boldsymbol{v}] = E[\boldsymbol{U_i} | \boldsymbol{v'}],$$

for any choices of $\boldsymbol{v}$, $\boldsymbol{v'}$. Therefore, agent i is indeed indifferent between any choice of $\boldsymbol{v}$. □

3.3 Discussion and Limits of the Approach

Our assumption of centrality is reasonable from a modelling perspective: Designing a system (which in this section is $\mathcal{M}[X \cup U]$ together with incoming nodes to this set) that allows for agent communication not through channels of the system is harder to model and often unreasonable. For example, in an anonymous internet application, communication beyond the one offered on websites difficult for the agents.

However, our definition of compatibility rules out an application to some of the results in the literature. We take [BEW]'s question as an example. They study whether for a single-agent MAID, the subdivision of an edge by a chance node can (for some parametrisation) increase utility. In particular, they wish to characterise for which MAIDs this is not possible for any choice of CPD on this new node. It is easy to see this is equivalent to the question whether a MAID and one where an additional *action* node is added subdividing an edge are equivalent.[3] Formally:

[3] As a deterministic action node implementing the identity function will yield the same expected utility for all agents.

Definition. *Let $\mathcal{M} = (X, D, U, E)$ be an influence graph and $e = (u, v) \in E$. Then we define the e-augmentation of $\mathcal{M}$ as the influence graph $\mathcal{M}^e := (X \cup \{d\}, D, U, E \setminus \{e\} \cup \{(u, d), (d, v)\})$ and assume that the domains of u and e are the same.*

Our algorithm helps us to find examples of MAIDs where an augmentation cannot change utility. The following is an example not contained in the cases listed in [BEW, Proposition 6]:

Proposition 2 (Irrelevance of node addition)**.** *Let $\mathcal{M}$ be an influence graph. $\mathcal{M}$ and its (u, v)-augmentation are equivalent if the following holds:* $\mathrm{Pa}(u) = \{r\} \subset D$ *and* $\mathrm{Ch}(u) \setminus \{v\} \subseteq D$ *and* $\mathrm{Pa}(r) \subseteq D$*, compare Figure 4.*

Proof. We can use line 2 of algorithm 1 in the non-augmented graph at node r and in the augmented one at u yielding the same graph. As these are equivalence transformations, $\mathcal{M}$ and its (u, v)-augmentation are equivalent. □

Our ground set is not rich enough to formulate some conditions for this question of equivalence. Indeed, it is not too hard to see that $u \in D$ and d-sep $(u, \mathrm{Desc}(v) \cap U | w)$ for any $w \in \mathrm{Ch}(u) \setminus \{v\}$ also implies that utility cannot be changed (see [BEW]). Unfortunately, the deletion of an incoming edge to a chance node is ruled out by our definition of compatibility. Therefore, one might think about a definition that also allows chance nodes to ignore some inputs. We leave this for further work.

Although limited in its theoretical guarantees, our standardising algorithm allows for simplifications that make equality testing easier and algorithms faster. We continue our study by considering a strategic choice of the nodes in X in the next two sections.

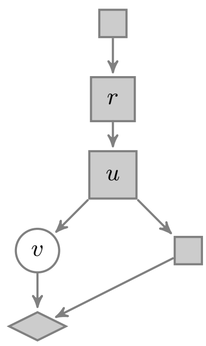

Figure 4: MAID for which (u, v)-augmentation cannot change utility.

4 Computing Interventions: Reducing to Non-Commitment

Where the last section studied canonical form of interactions with some nodes fixed, this section considers *strategic* choice of chance node CPDs. In this and the next section, we will study the following problem:

MAID COMMITMENT EQUILIBRIUM	
Input	A parametrised MAID $\mathcal{M}$ represented by lists of real numbers $(\widetilde{u}(\mathbf{Pa}(\boldsymbol{u})))_{\mathrm{Pa}(\boldsymbol{u})\in\mathrm{dom}(\mathbf{Pa}(\boldsymbol{u}))}$, $u \in U$, and $(p(\boldsymbol{x}\vert\mathbf{Pa}(\boldsymbol{x})))_{\substack{\mathbf{Pa}(\boldsymbol{x})\in\mathrm{dom}(\mathbf{Pa}(\boldsymbol{x}))\\ \boldsymbol{x}\in\mathrm{dom}(\boldsymbol{x})}}$, $x \in X$ (domains for $U \cup X \cup \mathrm{Pa}(U) \cup \mathrm{Pa}(X)$ are implicitly represented in the lists), A set $D^c \subseteq D_1$ represented by indices
Output	Decision rules $\boldsymbol{D^c}$ represented by lists of real numbers $(p(\boldsymbol{d}\vert\mathbf{Pa}(\boldsymbol{d})))_{\substack{\mathbf{Pa}(\boldsymbol{d})\in\mathrm{dom}(\mathbf{Pa}(\boldsymbol{d}))\\ \boldsymbol{d}\in\mathrm{dom}(\boldsymbol{d})}}$, $d \in D$ (domains for $D \setminus (U \cup X \cup \mathrm{Pa}(U) \cup \mathrm{Pa}(X))$ implicitly represented in the lists)
Goal	Choose $\boldsymbol{D^c}$ utility maximising among all decision rules $\boldsymbol{D^c}$ such that $\boldsymbol{D_1} \setminus \boldsymbol{D^c}$, $\boldsymbol{D_i}$, $i \geq 2$ are best responses.

It might not be obvious why this problem is different from MAID COMMITMENT EQUILIBRIUM. Reconsider the MAID from Figure 2 (reprinted for convenience as Figure 5) with the difference that b is now a commitment node (we denote commitment nodes by a double-lined square). Consider the parametrisation $\mathrm{dom}(\boldsymbol{u_a}) = \mathrm{dom}(\boldsymbol{u_b}) = \mathrm{dom}(\boldsymbol{a}) = \mathrm{dom}(\boldsymbol{b}) = \{0, 1\}$, $\widetilde{u_a}(\boldsymbol{a}, \boldsymbol{b}) = \delta_{\boldsymbol{ab}}\delta_{\boldsymbol{a}0}$ and $\widetilde{u_b}(\boldsymbol{a}, \boldsymbol{b}) = \delta_{\boldsymbol{ab}}\delta_{\boldsymbol{b}1}$, i.e. a would like both players to choose 0, b would like both players choose 1. In the equilibrium concept of MAID COMMITMENT EQUILIBRIUM, a would play 0 for sure and b would follow, in MAID EQUILIBRIUM, b can choose among any of her decision rules given a plays a best response to maximise her utility. The optimal choice for b is $p(\boldsymbol{b}\vert\boldsymbol{a}) = \delta_{\boldsymbol{b}1}$ and the optimal choice for a is $p(\boldsymbol{a}) = \delta_{\boldsymbol{a}1}$. Therefore, the „power“ in this parametrised MAID shifts due to commitment.

In terms of constraints, the difference between MAID COMMITMENT EQUILIBRIUM and MAID EQUILIBRIUM is the lack of a utility maximisation constraint for agent i. Denoting $\boldsymbol{(D \setminus D^c)_{D^c}}$ best responses to D^c, one has for MAID EQUILIBRIUM

$$\mathrm{EU}_i[(\boldsymbol{D^c}, \boldsymbol{(D \setminus D^c)_{D^c}})] \geq \mathrm{EU}_i[(\boldsymbol{D^{c\prime}}, \boldsymbol{(D \setminus D^c)_{D^{c\prime}}})] \tag{2}$$

instead of

$$\mathrm{EU}_i[(\boldsymbol{D^c}, \boldsymbol{(D \setminus D^c)_{D^c}})] \geq \mathrm{EU}_i[(\boldsymbol{D^{c\prime}}, \boldsymbol{(D \setminus D^c)_{D^c}})] \tag{3}$$

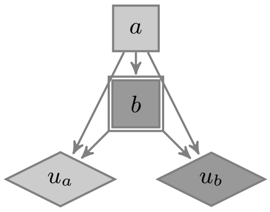

Figure 5: The strategic coordination MAID with commitment.

for MAID Commitment Equilibrium. We will call a choice of D^c satisfying Equation 2 a *commitment best response*.

We highlight that for single-agent settings, MAID Equilibrium is the same as MAID Commitment Equilibrium. Indeed, there is no difference between expected utility maximising choices and choices being best responses in a cooperative setting.

Another question is why we restrict ourselves to the case where $D^c \subseteq D_1$, i.e. all commitment nodes belong to one agent. We conclude this section with an example of an influence graph that has a parametrisation not admitting for a non-cooperative choice of commitment decision rules.

4.1 Finite Algorithm

A naïve algorithm for MAID Equilibrium might not even be finite, as there is a continuum of choices for the decision rules (they are probability distributions) and for each choice, the expected utility value would require an equilibrium computation.

In this section we present a first finite-time algorithm for the problem (assuming that the communication domains have bounded cardinality; we will study the size of communication domains in section 5). It will work in complete generality. As a downside, we show that not only does the reduction increase instance size, but also that a divide-and-conquer decomposition to speed up equilibrium computation as given in [KM03] is not possible.

Theorem 3 (Reduction to MAID equilibrium)**.** MAID Equilibrium $(\mathcal{M}, D^c)$ *can be reduced to* MAID Commitment Equilibrium $(\mathcal{M}')$.

The instance size transforms as follows: For $l := |D(\mathcal{M}) \setminus D^c|$, $m := |E(\mathcal{M})|$ *and* $k_d := |\mathrm{dom}(\mathbf{Pa}(\boldsymbol{d}))|$, $|V(M')| = n + 2\sum_{d \in D^c} k_d$ *and* $|E(M')| = m + \sum_{d \in D^c} 2k_d - 1 + (|\mathrm{Pa}(d)| - 1)k_d$.

Corollary 4 (Finite-time algorithm for general commitment MAIDs)**.** *Under the assumption that communication domain cardinalities are bounded, there is a finite-time algorithm for* MAID Equilibrium.

Proof of Corollary 4. Any MAID can be represented as an extensive form game [KM03, Section 4] and solved with a generic algorithm for computing equilibria in extensive form games [MM96]. □

The crucial observation for the proof is the following:

Proposition 5 (Unnecessary commitment). *Let $\mathcal{M}$ be an influence graph and D^c such that all nodes in D^c are roots. Then the set of* MAID *commitment equilibria and* MAID *equilibria with respect to D^c coincide.*

Proof. We have to prove that in this case, (2) and (3) coincide. Note that we can write for any agent

$$\mathrm{EU}_i\left[(\boldsymbol{D^c}, \boldsymbol{D} \setminus \boldsymbol{D^c})\right] = \sum_{\boldsymbol{D^c} \in \mathrm{dom}(\boldsymbol{D^c})} p(\boldsymbol{D^c}) \mathbb{E}\left[\sum_{u \in U_i} \widetilde{u}(\mathbf{Pa}(\boldsymbol{u})) \middle| \boldsymbol{D^c}\right].$$

Hence, the optimisation objective for each agent is a linear function of

$$\mathbb{E}\left[\sum_{u \in U_i} \widetilde{u}(\mathbf{Pa}(\boldsymbol{u})) \middle| \boldsymbol{D^c}\right].$$

Therefore, for each $\boldsymbol{D^c} \in \mathrm{dom}(\boldsymbol{D^c})$, the agents will separately maximise this value. Hence, for each deterministic, and by linearity of expectation also for any probabilistic, choice of $\boldsymbol{D^c}$, the other agents will only play best responses. Therefore, (3) and (2) coincide. □

Proof of Theorem 3. We use the same transformation as [MCZ12, Transformation 6]—which has been used to transform single-player influence diagrams keeping expected utility constant. We complement their result that their transformation in a multi-agent setting is a reduction for commitment equilibria.

For each commitment node $d \in D^c$ repeat the following. For all root commitment nodes, return the graph unchanged. Hence assume that $|\delta^-(d)| > 0$. For ease of notation, we will suppress the subscript of k_d. Let $\{\boldsymbol{P_i}\}_{i=1}^k = \mathrm{dom}(\mathbf{Pa}(\boldsymbol{d}))$ be an enumeration of all parent instantiations of d. Add k new action nodes c_i, $i \in [k]$ and k new chance nodes x_i, $i \in [k]$, each with $\mathrm{dom}(\boldsymbol{d})$. Add new edges from any parent of d to any x_i, $i \in [k]$. Furthermore, add edges (c_i, x_i), $i \in [k]$ and edges to make $x_1 \to x_2 \to \cdots \to x_k$ a path. Furthermore, connect x_k to all children of d. Compare Figure 6 for the construction. Having introduced new chance variables $\boldsymbol{x_i}$, $i \in [k]$, we need to define their parametrisation. For the first,

$$p(\boldsymbol{x_1} | \boldsymbol{c_1}, \mathbf{Pa}(\boldsymbol{d})) = \begin{cases} \delta_{\boldsymbol{P_1} \boldsymbol{x_1}} & \mathbf{Pa}(\boldsymbol{d}) = \boldsymbol{P_1} \\ \frac{1}{\mathrm{dom}(\boldsymbol{d})} & \text{else.} \end{cases}$$

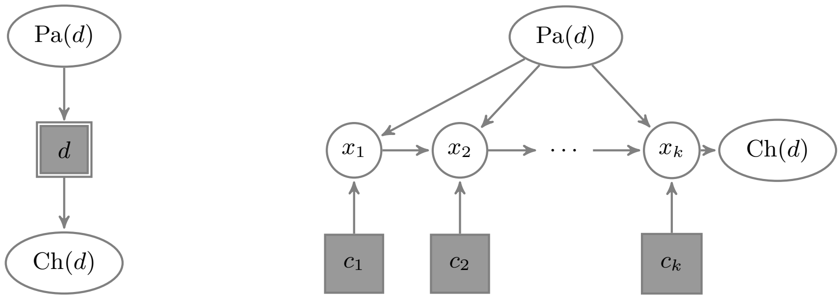

(a) MAID with commitment. (b) Reduction to MAID.

Figure 6: Transformation for the proof of Theorem 3.

(we remark that the denominator m in the corresponding definition in [MCZ12, Transformation 6] is wrong as this does not give a normalised distribution). For $i > 1$,

$$p(\boldsymbol{x_i}|\boldsymbol{c_i}\mathbf{Pa}(\widetilde{\boldsymbol{d}}), \boldsymbol{x_{i-1}}) = \begin{cases} \delta_{\boldsymbol{x_i}\boldsymbol{c_i}} & \mathbf{Pa}(\boldsymbol{d}) = \boldsymbol{p_i} \\ \delta_{\boldsymbol{x_i}\boldsymbol{x_{i-1}}} & \text{else.} \end{cases}$$

These conditional distributions for each fixed parent instantiation propagate a uniform random state until i is reached such that $\mathbf{Pa}(\boldsymbol{d}) = \boldsymbol{P_i}$. From then onward, the value in $\boldsymbol{c_i}$ is propagated and finally transmitted to $\mathrm{Ch}(d)$. One can think of the nodes c_i, $i \in [k]$ as an encoding the decision rule $\boldsymbol{d}$. For each parent instantiation, the random variables $\boldsymbol{x_i}$, $i \in [k]$ evaluate this function at the parent instantiations *on behalf of player i*.

The encoding of the decision rules is given by the following correspondence:

$$p(\boldsymbol{c_i}) = p_{\boldsymbol{d}|\mathbf{Pa}(\boldsymbol{D})}(\boldsymbol{d}|\boldsymbol{P_i}).$$

In [MCZ12, Proposition 7] it is shown for this correspondence and for fixed parametrisation and decision rules for $D \setminus \{d\}$ the same distribution on $V(\mathcal{M}) \setminus \{d\}$ is induced. In particular, other agents' best responses are the same before and after the transformation. Furthermore, the equality of distributions implies that the equilibrium utility by the two strategies for agent i are the same. It remains to show, that an optimal (commitment) choice for $\boldsymbol{d}$ corresponds to a best response $(c_1, c_2, \ldots, c_k)$ to the other agents. But note that all c_i, $i \geq k$ are roots. Therefore any best response is also a commitment best response by Proposition 5.

Concerning the number of nodes, there are exactly $2k$ additional nodes c_i and x_i. Furthermore, there are $2k - 1$ connecting new nodes among one another $(|\mathrm{Pa}(d)| - 1)k$ new edges connecting d's parents to the new nodes. Summing over all nodes $d \in D^c$ yields the claim. □

The coefficient k is doubly exponential in the number of inputs to the mechanism. We argue that this transformation has also been used in the computational literature [MCZ12].

Nevertheless, decomposing the problem into more manageable ones would be helpful. The next section shows that this is not the case by means of [KM03, Algorithm 6.2].

4.2 Negative Result for a Divide-and-Conquer Approach

[KM03, Algorithm 6.2] allows for faster equilibrium computation if the MAID has small strongly connected components. We show that for a fixed transformed commitment node $d \in D^c$, $(c_i)_{i\in[k]}$ is strongly connected in $\mathcal{M}'$'s relevance graph. This shows that a divide-and-conquer approach to equilibrium computation is not promising for MAID EQUILIBRIUM.

Proposition 6 (Large connected components in transformed Graph)**.** *Let $d \in D^c$ a non-root node and* $\mathrm{Desc}(d) \cap U_1 \neq \emptyset$ *in the un-transformed graph. Then the new nodes c_i introduced by the transformation in Theorem 3 pairwise strategically rely one another. In addition:*

1. *All c_i are relevant to or strategically rely on the same other nodes.*
2. *There are* MAIDs $\mathcal{M}$ *such that, in the transformed influence graph, the c_i strategically rely on nodes that d did not rely on in the original influence graph.*

Less technically, Proposition 6 says that commitment nodes for different parent instantiations are not selected separately if in the original graph d was relevant to the commitment agent's utility. This is consistent with the idea of „credible threats", i.e. that a player induces a desirable equilibrium by threatening agents with the consequences if they do not obey, even if this harms the player's own utility.

Proof of Proposition 6. We start with the main statement. Let $i, j \in [k]$. We use [KM03]'s graphical criterion for strategic relevance to show that c_i strategically relies on c_j. By $\mathrm{Desc}(d) \cap U_1 \neq \emptyset$, there is a directed path $x_{k_d} \dashrightarrow u_1$ from $x_{\tilde{k}}$ to U_1. For $i < j$, add a new parent $\hat{c}$ to c_j. Then

$$\hat{c} \to c_j \to x_j \to x_{j+1} \to \cdots \to x_{k_d} \dashrightarrow u_1$$

is an active path that establishes strategic relevance. For $i > j$, add a new parent $\hat{c}$, this time to c_i. Furthermore let d^- be a parent of d in the original graph. Then

$$\hat{c} \to c_i \to x_i \leftarrow d^- \to x_{k_d} \dashrightarrow u_1$$

is another active path that establishes strategic relevance in this case (note that the v-structure at x_i is unblocked by the evidence from c_j (which is a descendant of x_i by $i > j$).

We continue with the additional assertions. It is straightforward to check that the sets of descendants and the set of d-connected nodes to c_i (and, as they have no parents, also

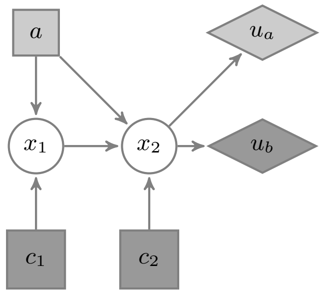

Figure 7: Strategic relevance through addition of commitment.

from a newly added parent $\hat{c}$) coincide for any $i \in [k_d]$. As no path can pass through c_i (it has only one incident node) and as it has no parents, the conditioning on c_i or its parents neither blocks nor unblock any active path. Therefore, all c_i are s-reachable and s-reach the same nodes.

For the last assertion, the influence graph Figure 5 is an example. The transformed influence graph is depicted in Figure 7. Note that in Figure 5, there was no active path from $\hat{a}$, a new parent added to a, to u_b given $b \cup \mathrm{Pa}(b)$ (as the only path is blocked by b), but in the transformed MAID, there is the active path

$$c_1 \to x_1 \to x_2 \to u_b,$$

yielding the claim. □

The number of additional nodes c_i far exceeds the number of nodes in the graph. Therefore the fact that they all form a strongly connected component limits the usefulness of the decomposition in [KM03, Theorem 6.2].

Therefore, algorithmic speedup due to divide-and-conquer is limited, and the strategic relevance of commitment nodes can be studied via the reduction given in Theorem 3.

We close this section by showing that a non-cooperative problem similar to MAID Equilibrium suffers from equilibrium non-existence.

4.3 The Case of Several Agents

Consider the modification of MAID Equilibrium for $D^c \cap D_i \neq 0$ and $D^c \cap D_j \neq \emptyset$ for some $i \neq j \in [n]$. Now the decision rules of different commitment agents are chosen so that they are best responses to one another given that all other decision nodes are best responses to them and each other. The following example appears in different formalism in [Mye81, Proposition 3]. We reformulate it as a problem of computing an MAID Equilibrium where two agents have commitment nodes. Therefore, even analytically, a non-cooperative version of MAID Equilibrium is challenging.

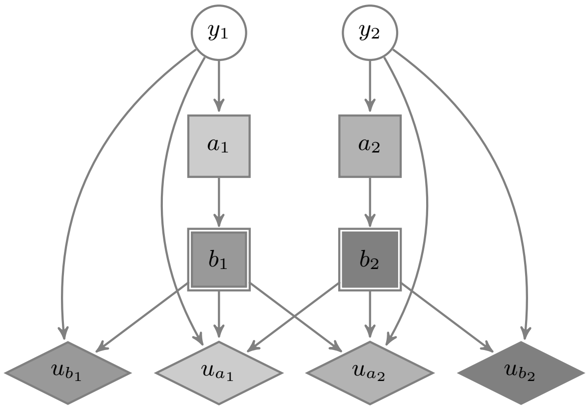

Figure 8: MAID not admitting a commitment equilibrium

Proposition 7 (Equilibrium non-existence)**.** *Figure 8 has a parametrisation such that there is no non-cooperative commitment equilibrium.*

Proof. We give the example of [Mye82] with one additional argument concerning communication domains. Parametrise

$$\mathrm{dom}(\boldsymbol{y_1}) = \mathrm{dom}(\boldsymbol{y_2}) = \{0, 1\} \qquad \mathrm{dom}(\boldsymbol{b_1}) = \mathrm{dom}(\boldsymbol{b_2}) = \{0, 1, 2\}$$

The utility functions $\widetilde{u_{a_i}}$, $\widetilde{u_{b_i}}$, $i = 1, 2$ are shown in Table 1 with implied domains in the real line. We first remark that $\mathrm{dom}(\boldsymbol{a_1})$ and $\mathrm{dom}(\boldsymbol{a_2})$ can be chosen as $\{0, 1\}$. Indeed,

$\widetilde{u_{a_1}}, \widetilde{u_{b_1}}$	$\boldsymbol{y_1} = 0$	$\boldsymbol{y_1} = 1$
$\boldsymbol{b_1} = 0$	$1, 6$	$2 - \delta_{\boldsymbol{b_2}2}, 0$
$\boldsymbol{b_1} = 1$	$2 - \delta_{\boldsymbol{b_2}2}, 0$	$1, 6$
$\boldsymbol{b_1} = 2$	$0, 5$	$0, 5$

$\widetilde{u_{a_2}}, \widetilde{u_{b_2}}$	$\boldsymbol{y_2} = 0$	$\boldsymbol{y_2} = 1$
$\boldsymbol{b_2} = 0$	$1, 6$	$1 + \delta_{\boldsymbol{b_1}2}, 0$
$\boldsymbol{b_2} = 1$	$1 + \delta_{\boldsymbol{b_1}2}, 0$	$1, 6$
$\boldsymbol{b_2} = 2$	$0, 5$	$0, 5$

Table 1: Utilities for the example of a MAID not admitting an equilibrium with several commitment agents.

for any other domains and functions, let $p_{\boldsymbol{a_i}|\boldsymbol{y_i}}(\boldsymbol{a_i}|\boldsymbol{y_i})$ and $p_{\boldsymbol{b_i}|\boldsymbol{a_i}}(\boldsymbol{b_i}|\boldsymbol{a_i})$, $\boldsymbol{a_i} \in \mathrm{dom}(\boldsymbol{a_i})$ be given. Then consider the conditional probability mass functions

$$p^{\frown}(\boldsymbol{a_i}|\boldsymbol{y_i}) = \delta_{\boldsymbol{y_i}} \qquad p^d(\boldsymbol{b_i}|\boldsymbol{a_i}) = \sum_{\boldsymbol{a_i'} \in \mathrm{dom}(\boldsymbol{a_i})} p_{\boldsymbol{b_i}|\boldsymbol{a_i}}(\boldsymbol{b_i}|\boldsymbol{a_i'}) p_{\boldsymbol{a_i}|\boldsymbol{y_i}}(\boldsymbol{a_i'}|\boldsymbol{a_i})$$

It is straightforward to check that $p^{\frown}$ and p^d induce the same distribution on u_{a_i} and u_{b_i} as $p_{\boldsymbol{a_i}|\boldsymbol{y_i}}$ and $p_{\boldsymbol{b_i}|\boldsymbol{a_i}}$. Therefore, there are decision rules that give each agent the same utility. As any decision rule can be transformed into such a standard rule (we will study $p^{\frown}$, so-called concatenation rules, in more detail in section 5) this even is an equilibrium.

Therefore, we can restrict ourselves to

$$\mathrm{dom}(\boldsymbol{a_i}) = \{0, 1\}$$

The following four implications for marginal distributions suffice to establish equilibrium non-existence

$$\begin{aligned} p_{\boldsymbol{b_2}}(2) = 1 &\Rightarrow p_{\boldsymbol{b_1}}(2) = 0 \\ p_{\boldsymbol{b_1}}(2) = 0 &\Rightarrow p_{\boldsymbol{b_2}}(2) = 0 \\ p_{\boldsymbol{b_2}}(2) < 1 &\Rightarrow p_{\boldsymbol{b_1}}(2) = 1 \\ p_{\boldsymbol{b_1}}(2) > 0 &\Rightarrow p_{\boldsymbol{b_2}}(2) = 1 \end{aligned}$$

First, observe that this is not satisfiable. In the case that $p_{\boldsymbol{b_2}}(2) = 1$, the first two implications contradict each other. In the case $p_{\boldsymbol{b_2}}(2) < 1$, the last two contradict each other.

If $p_{\boldsymbol{b_2}}(2) = 1$, then a_1 is indifferent between $\boldsymbol{b_1} \in \{0, 1\}$ but prefers it to $\boldsymbol{b_1} = 2$. b_1 prefers to match the signal $\boldsymbol{y_1}$. Therefore, in any equilibrium,

$$p(\boldsymbol{a_1}|\boldsymbol{y_1}) = \delta_{\boldsymbol{a_1}\boldsymbol{y_1}} \qquad\qquad p(\boldsymbol{b_1}|\boldsymbol{a_1}) = \delta_{\boldsymbol{b_1}\boldsymbol{a_1}},$$

in particular $p_{\boldsymbol{b_1}}(2) = 0$. A very similar argument on a_2 and b_2 holds for the second implication—we leave it to the reader.

In case $p_{\boldsymbol{b_2}}(2) < 1$, a_1 prefers if b_1 does not match the signal $\boldsymbol{y_1}$. Then, in any equilibrium, it is optimal for a_1 to not disclose any information on y_1 to b_1, upon which b_1 chooses 2,

$$p(\boldsymbol{a_1}|\boldsymbol{y_1}) = \frac{1}{2} \qquad\qquad p(\boldsymbol{b_1}|\boldsymbol{a_1}) = \delta_{\boldsymbol{b_1}2}.$$

This implies $p_{\boldsymbol{b_1}}(2) = 1$. The fourth implication is again very similar. □

It is not apparent why the reduction from Theorem 3—which would guarantee equilibrium existence—breaks down in the case of multiple equilibrium agents. The problem here, however, is that the construction of the reduction by making all commitment nodes roots frees commitment decision rules from *all* incentives. Therefore, using the same reduction for several agents would mean that these choose their decision rules cooperatively, which is often not reasonable. In cooperative applications, this could be, however, a potential direction for further research.

In this section, we presented a first approach to computing MAID commitment equilibria and showed its limits. The next section makes a further assumption to allow for an LP formulation which is more efficient.

5 Computing Interventions: Characterising Strategies and an LP Formulation

This section shows for a class of MAIDs with a restricted structure that any attainable expected utility vector is attained for one fixed choice of communication domains and decision rules only depending on the parametrisation's domains. Using this structure, we show for a narrowly defined class of MAIDs how polynomially solvable LP formulations for MAID Equilibrium may be obtained.

The outline of the section is as follows: We first define the class of *private values* MAIDs and then introduce the main result, a characterisation of the equilibrium decision rules for action nodes in $D \setminus D^c$. Then, we introduce our class of centralised liability MAIDs. We show that liability MAIDs admit a polynomially solvable LP formulation for MAID Equilibrium in the case of dense dependency.

Definition. *Let* $\mathcal{M} = (X, D, U, E)$, $D^c \subseteq D_1$ *be an instance of* MAID Equilibrium. *We call the tuple* $(\mathcal{M}, D^c)$ *a* private values MAID *if the following four properties hold:*

No information restrictions $\mathrm{Pa}(X) \cap D = \emptyset$.

No non-commitment utility-relevant actions $\mathrm{Ch}(v) \cap U \neq \emptyset$ *only if* $v \in D^c \cup X$.

Independent Actions *For any* $d, d' \in D^c$, $(\mathrm{Anc}(d) \cap \mathrm{Anc}(d') \cap (D \setminus D^c)) = \emptyset$ *or* $\mathrm{Pa}(d) = \mathrm{Pa}(d')$.

No Hints $E(D^c, D) = \emptyset$.

Compare Figure 9 for a pictorial description of private value MAIDs.

We comment on the assumptions in the definition. Allowing for communication restrictions in mechanism design is complicated, and only partial results are known, e.g. [BNS07]. Similarly, there are only few results when outcome-relevant actions by non-commitment nodes are allowed. Some of the resulting issues are covered in dynamic mechanism design; see survey [BV19]. The third assumption ensures that there is no communication between agents that send their messages to different commitment nodes. Being a tractability assumption, this might be interesting for modelling decentralised decision making. We leave the relaxation of the independent actions assumption for further work. The last assumption says that the commitment nodes do not send signals to agents, which says that information flows only in one direction and commitment nodes do not take coordinating roles for agent decisions.

We highlight that private values MAIDs allow for inter-agent communication.

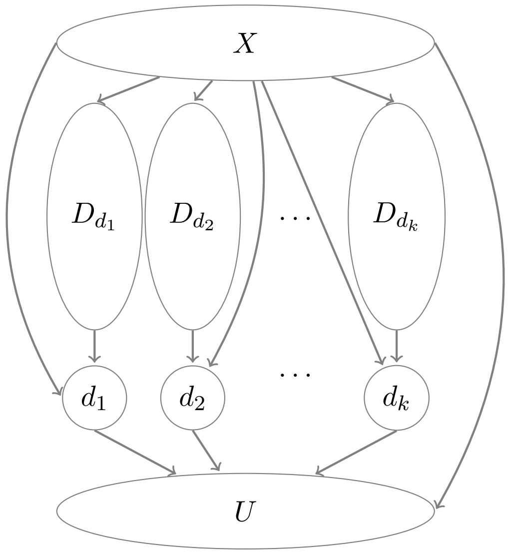

Figure 9: Structure of private value maids. The only edges in this structures graph are either within blocks or between the blocks connected by edges in this graph. For a definition of D_d, see the proof of Theorem 8.

5.1 Concatenation Strategies

Definition (Concatenation Strategies)**.** *Let $d \in D \setminus D^c$. The* concatenation strategy *consists of the communication domain specification* $\mathrm{dom}(\boldsymbol{d}) := \mathrm{dom}(\mathbf{Pa}(\boldsymbol{d}))$ *together with decision rule*

$$p^{\frown}(\boldsymbol{d}|\mathbf{Pa}(\boldsymbol{d})) = \delta_{\boldsymbol{d}\mathbf{Pa}(\boldsymbol{d})}$$

that deterministically sends all parent information.

Note that by our no information restrictions assumption, if all $d \in D \setminus D^c$ play concatenation strategies, one can construct domains for all agents in topological order of the DAG $(X \cup D \cup U, E)$. The following result shows that suffices to replicate all possible MAID equilibria.

Theorem 8 (Commitment strategies)**.** *Let $(\mathcal{M}, D^c)$ be a private values* MAID *and let $(\boldsymbol{D^c}, \boldsymbol{D} \setminus \boldsymbol{D^c})$ be a* MAID *equilibrium. Then there are decision rules $p^d(\boldsymbol{d}|\mathbf{Pa}(\boldsymbol{d}))$, $d \in D^c$, such that the induced joint distribution $\boldsymbol{D'}$ is a* MAID *equilibrium.*

Proof. As a first step, we can merge all nodes $d \in D^c$ that have the same parent set into one node that has their cartesian product as domain. This follows, because the parametrisation in which each utility node only takes into account the part that belonged to its parent in the original graph implies the same utilities. For this reason, it is without loss to assume that

$$(\mathrm{Anc}(d) \cap \mathrm{Anc}(d') \cap (D \setminus D^c)) = \emptyset$$

for any $d, d' \in D^c$.

For brevity, we denote in this proof $D^{nc} = D \setminus D^c$. Denote furthermore

$$X_d^c \coloneqq \mathrm{Pa}(d) \cap X \qquad X_d \coloneqq \mathrm{Anc}(d) \cap X \setminus X_d^c$$
$$D_d^c \coloneqq \mathrm{Pa}(d) \cap D^{nc} \qquad D_d \coloneqq \mathrm{Anc}(d) \cap D^{nc} \setminus D_d^c$$

all the values in X resp. D accessible to $d \in D^{nc}$ via concatenation strategies and those directly observed by d. Denote $\hat{X}_d \coloneqq X_d^c \cup X_d$ and $\hat{D}_d \coloneqq X_d^c \cup X_d$, and furthermore $\hat{X} = \bigcup_{d \in D^c} \hat{X}_d$ and $\hat{D} = \bigcup_{d \in D^c} \hat{D}_d$. For any subset $A \subseteq \hat{X}_d \cup \hat{D}_d$ denote by

$$\boldsymbol{A}(\mathbf{Pa}(\boldsymbol{d}))$$

the value $\boldsymbol{A}$ as either observed by d or sent via concatenation strategies from $d \in D_d$ (this value is unique as in equilibrium, all observed or via a concatenation strategy reported values must coincide).

With respect to $(\boldsymbol{D^c}, \boldsymbol{D^{nc}})$, let

$$p(\boldsymbol{D^{nc}}|\boldsymbol{X}) \qquad p(\boldsymbol{D^c}|\boldsymbol{D^{nc}} \cup \boldsymbol{X})$$

be conditional probability mass functions implied by set of decision rules $\boldsymbol{D}$. Denote the domains with respect to $(\boldsymbol{D^c}, \boldsymbol{D^{nc}})$ as dom (we will not need a notation for the domains of a concatenation strategy). We can simplify the probability mass functions through the following independencies, which can be checked via the d-separation criterion.

$$\begin{aligned}\hat{\boldsymbol{D}}_{\boldsymbol{d}} &\perp\!\!\!\perp \hat{\boldsymbol{D}}_{\boldsymbol{d'}} \quad |\hat{\boldsymbol{X}}_{\boldsymbol{d}} \\ \hat{\boldsymbol{D}}_{\boldsymbol{d}} &\perp\!\!\!\perp \boldsymbol{X} \setminus \hat{\boldsymbol{X}}_{\boldsymbol{d}}|\hat{\boldsymbol{X}}_{\boldsymbol{d}}.\end{aligned}$$

for $d, d' \in D^c$, $d \neq d$. Note furthermore, that we can assume that $\hat{D} = D^{nc}$ any node in the complement cannot have a directed path to a utility node. Then

$$\begin{aligned}p(\boldsymbol{D^{nc}}|\boldsymbol{X}) &= \prod_{d \in D^c} p(\hat{\boldsymbol{D}}_{\boldsymbol{d}}|\hat{\boldsymbol{X}}_d) \\ p(\boldsymbol{D^c}|\boldsymbol{D^{nc}} \cup \boldsymbol{X}) &= \prod_{d \in D^c} p(\boldsymbol{d}|\hat{\boldsymbol{D}}_{\boldsymbol{d}}).\end{aligned}$$

Then for $d \in D^c$, define

$$p^d(\boldsymbol{d}|\mathbf{Pa}(\boldsymbol{d})) = \sum_{\hat{\boldsymbol{D}}_{\boldsymbol{d}} \in \mathrm{dom}(\hat{\boldsymbol{D}}_{\boldsymbol{d}})} p(\boldsymbol{d}|\hat{\boldsymbol{D}}_{\boldsymbol{d}}(\mathbf{Pa}(\boldsymbol{d})))p(\hat{\boldsymbol{D}}_{\boldsymbol{d}}|\hat{\boldsymbol{X}}_{\boldsymbol{d}}(\mathbf{Pa}(\boldsymbol{d})))$$

It is tedious but straightforward to check that $p^\frown$ and p^d, $d \in D^c$ induce the same probability distribution on D as $\boldsymbol{D}$. This also implies that all expected utilities of agents are the same.

Note that as for any strategies this gives the agents the same utility, the decision rules defined by $(p^\frown, p^d)$ form an equilibrium. This concludes the proof. $\square$

Results similar to Theorem 8 are called *revelation principle* in the economics literature, compare [Mye86] for an early form. Our result differs from prior known revelation principle in that agents can report complexly interrelated private information in the form of chance nodes $\boldsymbol{X}$. Requiring more structure of the MAIDs and on the parametrisation, this characterisation can even lead to a compact LP solution, as the next subsection shows.

5.2 Using Properties of a Parametrisation for Efficient Computation

Given the characterisation Theorem 8, a linear programming formulation of MAID Equilibrium is possible. In complex communication patterns, however, this formulation is vast, despite finite.

We restrict ourselves to a very particular structure that allows for pre-processing that yields a polynomially solvable LP.

Definition. *A liability* MAID *is a parametrised private values* MAID $(\mathcal{M}, D^c)$ *such that the following holds:*

Centralisation $E(D_i, D_j) = \emptyset$, $i \neq j \in [n]$, *compare definition in section 3.*

One Commitment Node Except Liability Nodes $D^c = \{d^c\} \cup \{d^l_j\}_{j\in[n]\setminus\{1\}}$.

Unlimited Liability *For each* $i \in [n] \setminus \{1\}$, *there is* $u^l_i \in U_i$, $\mathrm{Pa}(u^l_i) = \mathrm{Pa}(d^c)$. *Furthermore,* $\max \widetilde{u^l_i} - \min \widetilde{u^l_i} > \sum_{u\in U_i\setminus\{u^l_i\}} \max \widetilde{u} - \min \widetilde{u}$.

We say that for a sequence of parametrised MAIDs $(\mathcal{M}_n)_{n\in\mathbb{N}}$ correlation is dense *if* $\sum_{v\in D(\mathcal{M}_n)} |\delta^-(v)| \in \Theta(|X(\mathcal{M}_n|)$. *This is in particular satisfied if one node has in-degree in* $\Theta(|X(\mathcal{M}_n)|)$.

The first assumption captures the property that all communication is controlled by the commitment nodes, the second and third say that for each non-commitment agent there is one commitment node that is exactly utility relevant for one agent and that the corresponding utility node has a large range. The name *liability* will become clear in the proof of the next proposition.

Proposition 9 (Polynomial-time algorithm for liability MAIDs). *There is an* LP *for* MAID Equilibrium *liability* MAIDs *that can be solved in polynomial time if there is a uniform bound C on the cardinalities of domains* $\mathrm{dom}(\boldsymbol{x})$, $x \in X$.

Proof. By Theorem 8 (ancestor sets of all commitment nodes are equal), we can assume that all choice rules except for $\boldsymbol{d}$, $\boldsymbol{d_i^l}$, $i \in [n] \setminus \{1\}$ are concatenation rules. We transform the instance of MAID Equilibrium to a problem of determining only one choice rule for $\boldsymbol{d^c}$. We need additional notation: For $d \in \mathrm{Pa}(d^c) = \mathrm{Pa}(d_j^l)$, $j \in [n] \setminus \{1\}$, $d \in D_j$ and $x \in \mathrm{Anc}(d^c) \cap X$, denote $\boldsymbol{d^{(x)}}$ the report of $\boldsymbol{x}$ by d. If this cannot be defined (as d reports different values of x), let u_j^l take its minimal value with certainty. Otherwise, if $\boldsymbol{d^{(x)}} \neq \boldsymbol{d'^{(x)}}$ for $d \in D_i$, $d' \in D_j$, $i, j \in [n] \setminus \{1\}$, then give u_i^l and u_j^l their minimal values. In any other case, give an agent the maximal value of u_j^l. Note that this (weakly) increases the utility of playing a concatenation strategy and (weakly) decreases the utility of not playing a concatenation strategy. Therefore, also after this change of choice rule, concatenation strategies form an equilibrium. Furthermore, this does not change the utility of agent 1. Therefore, this is still a MAID commitment equilibrium. Under these decision rules for d_i^l, furthermore, agents $j \in [n] \setminus \{1\}$ would prefer any decision at $\boldsymbol{d}$ over the minimum value at u_j^l, which implies that they will never misreport a value that is reported by at least one other agent (and will themselves not report different values). Note that at this point the assumption of centrality is crucial: If on a path from X to d^c there were different agent only giving the last agent negative utility would not suffice to let agent report truthfully.

Therefore, as an equivalent problem, we can assume that $D^c = \{d^c\}$ and ignore all utility nodes u_j^l, $j \in [n] \setminus \{1\}$. We can define (differently as in the proof of Theorem 8)

$$X_i = \left(\bigcup_{d \in D_i} \mathrm{Anc}(d) \cap X \right) \setminus \left(\bigcup_{d \in D \setminus D_i} \mathrm{Anc}(d) \cap X \right)$$

to be the set of all chance node values that are uniquely reported by agent $i \in [n] \setminus \{1\}$ and not observed by d^c. Furthermore, denote

$$\hat{X} = \mathrm{Anc}(D^c)$$

the values that are either observed by d^c or reported by an agent (but maybe multiple times). Furthermore, we will write a conditional probability mass function $p(\boldsymbol{d}|\boldsymbol{X})$ in sequence notation $p_{\boldsymbol{d}|\boldsymbol{X}}$. In addition, we write $u^i_{\boldsymbol{d},\hat{\boldsymbol{X}}}$ instead of

$$\sum_{u \in U_i} \mathbb{E}[\widetilde{u}(\boldsymbol{d}, \hat{\boldsymbol{X}}, \boldsymbol{X} \setminus \hat{\boldsymbol{X}}) | \boldsymbol{d}, \hat{\boldsymbol{X}}].$$

Furthermore, we write d instead of d^c. With these notations, the following LP solves the

MAID commitment equilibrium problem

$$\text{maximise } p^T u = \sum_{\substack{\hat{\boldsymbol{X}} \in \operatorname{dom}(\hat{\boldsymbol{X}}) \\ \boldsymbol{d} \in \operatorname{dom}(\boldsymbol{d})}} p_{\hat{\boldsymbol{X}}} p_{\boldsymbol{d}|\hat{\boldsymbol{X}}} u_{\boldsymbol{d},\hat{\boldsymbol{X}}}, \qquad \text{such that}$$

$$\sum_{\substack{\hat{\boldsymbol{X}} \setminus \boldsymbol{X_j} \in \operatorname{dom}(\hat{\boldsymbol{X}} \setminus \boldsymbol{X_j}) \\ \boldsymbol{d} \in \operatorname{dom}(\boldsymbol{d})}} p_{\hat{\boldsymbol{X}}} p_{\boldsymbol{d}|\hat{\boldsymbol{X}} \setminus \boldsymbol{X_j}, \boldsymbol{X_j}} u^i_{\boldsymbol{d},\hat{\boldsymbol{X}} \setminus \boldsymbol{X_j}, \boldsymbol{X_j}} \qquad i \in [n] \setminus \{1\},$$

$$\geq \sum_{\substack{\hat{\boldsymbol{X}} \setminus \boldsymbol{X_j} \in \operatorname{dom}(\hat{\boldsymbol{X}} \setminus \boldsymbol{X_j}) \\ \boldsymbol{d} \in \operatorname{dom}(\boldsymbol{d})}} p_{\hat{\boldsymbol{X}}} p_{\boldsymbol{d}|\hat{\boldsymbol{X}} \setminus \boldsymbol{X_j}, \boldsymbol{X}'_{\boldsymbol{j}}} u^i_{\boldsymbol{d},\hat{\boldsymbol{X}} \setminus \boldsymbol{X_j}, \boldsymbol{X}'_{\boldsymbol{j}}} \qquad \boldsymbol{X_j} \in \operatorname{dom}(\boldsymbol{X_j})$$

$$\sum_{\boldsymbol{d} \in \mathbf{dom}\,(\boldsymbol{d})} p_{\boldsymbol{d}|\operatorname{dom}(\hat{\boldsymbol{X}})} = 1, \qquad \hat{\boldsymbol{X}} \in \operatorname{dom}(\hat{\boldsymbol{X}}),$$

$$p_{\boldsymbol{d}|\operatorname{dom}(\hat{\boldsymbol{X}})} \geq 0, \qquad \boldsymbol{d} \in \operatorname{dom}(\boldsymbol{d})$$

$$\hat{\boldsymbol{X}} \in \operatorname{dom}(\hat{\boldsymbol{X}}).$$

Here, the vector p is

$$(p_{\boldsymbol{d}|\hat{\boldsymbol{X}}})_{\substack{\boldsymbol{d} \in \operatorname{dom}(\boldsymbol{d}) \\ \hat{\boldsymbol{X}} \in \operatorname{dom}(\hat{\boldsymbol{X}}}}$$

The last two sets of constraints ensure that $p_{\boldsymbol{d}|\operatorname{dom}(\hat{\boldsymbol{X}})}$ indeed defines a probability distribution. The objective maximises expected utility for the commitment agent. The first set of constraints ensures that truth-telling is a best response for the values of X he does not report himself.

It remains to show that this formulation is polynomial-time solvable. First observe that $p_{\hat{\boldsymbol{X}}}$ can be computed using variable elimination in time $O(|V(M)| \cdot |\operatorname{dom}(\boldsymbol{X})|)$ by variable elimination. Furthermore, the sum of expected utilities $u_{\boldsymbol{d},\hat{\boldsymbol{X}}}$ can also be computed using variable elimination in time $O(|U(M)| \cdot \operatorname{dom}(\boldsymbol{X}))$, compare [KFB09]. $(|\operatorname{dom}(d)| + 2) \cdot |\operatorname{dom}(\boldsymbol{X})|$ is an upper bound on the number of inequalities and $|\operatorname{dom}(\hat{\boldsymbol{X}})| \cdot |\operatorname{dom}(\boldsymbol{d})|$ is the number of variables. Both are polynomial in the representation size of the parametrisation, as for sufficiently large n, there must be a node $\boldsymbol{X}$ that has in-degree at least $c|X(\mathcal{M}_n)|$ and hence needs at least $2^{c|X(\mathcal{M}_n)|}$ representational size (as a preprocessing step, we can just ignore all chance variables with domain cardinality one, as they are deterministic). Furthermore, the utility functions need a representation size of at least $|\operatorname{dom}(\boldsymbol{d})|$. Finally,

$$|\operatorname{dom}(\boldsymbol{X})| \leq C^{|X(\mathcal{M}_n)|} \leq (2^{c|X(\mathcal{M}_n)|})^{\frac{1}{c}\log_2(\frac{C}{2})}$$

proving that the number of variables and the equations is polynomial in the representation size if $\mathcal{M}_n$. Therefore, the separation problem can be solved in polynomial time, implying a polynomial-time algorithm by the ellipsoid method (see e.g. [GLS93]). □

This section showed that in private values settings, the communication domains and decision rules can be chosen in a particular form. Together with an assumption on parametrisation, this can be turned into an efficiently solvable LP for MAID Equilibrium.

We presented a second approach to solving MAID Equilibrium, which might be promising for subclasses of MAIDs, as we demonstrated. We now present related work and, consequently, conclude.

6 Related Work

Our first contribution on canonical forms of games is most similar to the literature on the *reduced form* of extensive-form games. [Tho52] defined four transformations as „equivalence transformations“. They showed that for deterministic games, these transformations allow for a unique minimal game in a sense made precise in the working paper. [KM86] extend this result to games including chance (see also the unified framework [Bru]) and uses it in a classification of equilibrium concepts. [LF87] obtains similar results with a focus on identifying strategies of agents. We differ by our more unified approach. First, our model (MAID) allow for a richer temporal structure. Furthermore, the game transformations mentioned above are not found by defining an equivalence relation but are *ad-hoc* choices.

A second part of the literature that our first contribution relates to is system design. There has been work on insertion of in-edges to decision nodes, among others [LN01; Mil+08; Eve+19]. All give either partial (e.g. single-action) or only one direction (namely that d-separation in the graph implies that any equilibrium in the graph without the edge also an equilibrium when adding the edge). [Bla53; How66; BEW] study the addition of nodes to influence diagrams, [Bla53] in a restricted setting, [How66] proposing an algorithm, [BEW] with sufficient criteria that the utility of agents is not changed by the introduction of new nodes. Furthermore, there has been work on whether the control of a chance node by an agent can change utilities for agents [Eve+19; SH10]. Our approach does not consider specific operations on graphs and check whether they do not change equilibrium outcomes, but define a concept of equivalence. By providing a normalising algorithm, we unify several operations that do not change outcomes.

Our second and third contributions are related to algorithmic mechanism design (which in influence graph terms would be „decision rule design“). On the one hand, algorithmic mechanism design that uses the equilibrium concept of dominant-incentive compatible equilibria (hence all actions should be best responses whatever the other agents' decision rules are) studies (combinatorial) auctions in different informational and computing settings. See the survey [Nis15] for an introduction to this field. Our approach differs in that randomness plays a vital role in our analysis, for which a more permissive equilibrium concept is needed. Furthermore, we do not require combinatorial structure in the problem such as valuations on subsets of a finite set, but only assume that all our state variables have finite domains.

Bayesian algorithmic mechanism design takes a perspective of approximation when the environment is not exactly known to the designer. Indeed, as [Har13] (see also the chapter in [Nis+07]) shows, decision rules can adapt to information they get from a subset of the agents without raising incentive problems. Using such algorithms, constant-factor approximations on welfare of interventions can be obtained. We differ from this literature in that we do assume that all agents have *complete* information of probability distributions. The hardness in our problem arises from the complicated correlation structure of information and not from the interactions of learning model parameters and incentives as in [Har13].

Our transformation in section 4 introduces much communication into a game theoretic model. Closest to this explicit modelling of communication of the specifics of a mechanism is [NS06]. They study a combinatorial allocation model (i.e. there are several goods which shall be partitioned into sets allocated to different agents, which then pay pre-specified tranfers). In their model, they show an exponential lower bound on communication whenever the auctioneer would like more expected surplus than when allocating all items in grand bundle. In particular it is shown that almost all personalised prices must be communicated to the agents. Our model is different in that we do assume neither combinatorial structure nor specific functional forms as [NS06] does. We are however similar in that in our transformation, the complete mechanism is communicated to chance nodes.

The non-existence result we presented in section 4 can be relieved by a concept of quasi-equilibrium, as shown in [Mye82, Section 4]. In other models where interactions between agents and commitment choice nodes are required to be denser, equilibria exist, as [Yam10] shows in their work on several competing principals (i.e. commitment choice nodes of different agents). We differ from this in that we neither relax the equilibrium assumptions imposed on MAIDs nor consider restricted graph classes, raising a non-existence issue.

Our last contribution on a characterisation of communication domains and, implied, a LP reduction of MAID EQUILIBRIUM is in the tradition of the revelation principle [Mye82]. The literature has produced variants with dynamic arrival of information, such as in [Mye86] and for more different equilibrium concepts [SW17]. For a model with, in out terminology, three players, one commitment node and two chance nodes, [CS02] gave an LP reduction for MAID EQUILIBRIUM. [CS04] showed that concatenation strategies in a restricted setting could make the computational problem much harder for the agent possessing commitment nodes. We differ from the first three papers in that our notation is lighter than theirs and in that we allow for sparsity in the model. We consider a much more general model than the fourth paper in deriving our results. In addition, the complexity results in [CS04] (in particular their Theorem 2) relate to deterministic mechanisms, which much more often produce computational hardness result. It is not evident whether this result transfers to a Bayesian setting.

Finally, Algorithmic Game Theory studied other graphical representations of games. Strategic graph formation [Jac10, chapter 11] and congestion/routing games [Nis+07, chapter 18] take the graph structure as part of the game. Local/graphical games [Nis+07, chapter 7] let only agents interact that are incident in an underlying graph. Action-graph games [JLB11] allow the actions to only depend on the action counts of some other agents, hence have sparsity enforced depending on context. Temporal models such as extensive form game representations or the similar concept of decision trees are most similar to influence diagrams in nature (compare [MWG95, chapter 9] resp. [KFB09, Section 23.1] for introductions from an economic or probabilistic graphical models perspective). By using the established MAID formalism, we allow for some temporal structure, but not a total order, which none of the other models permits.

Having given pointers to relevant literature, we are ready to conclude.

7 Conclusion and Future Directions

This thesis asked for a concise model for interventions on agent interactions and how to compute such interventions. We obtained some first results, which are the following:

First, we defined an equivalence relation on Multi-Agent influence diagrams (*outcome equivalence*) and presented a normalising algorithm. This allowed us to identify a slight generalisation of results in [BEW].

Furthermore, we showed that all MAIDs with commitment may be represented using ordinary MAIDs, but that this reduction adds exponentially many nodes. We also showed that an approach for dividing a MAID computation into smaller sub-problems does not help, as nodes newly introduced in the reduction have strong dependencies.

Finally, we showed that for a restriction of agent communication, the strategies of agents can be chosen in a way that allows for the solution of the problem as an LP.

We highlight three shortcomings of the present work.

On the one hand, we did not consider the design of channels in case there are limits to communication. Indeed, in section 3, we treated all chance nodes as fixed, and in section 5 we assumed that are no chance nodes on any path connecting two decision nodes.[4]

Furthermore, our formalism does not allow us to study specific parametrisations such as functional forms for utilities. These, however, are important for many of the standard results in the (economic) mechanism design literature. For example, [Mye81]'s celebrated results depend crucially on quasi-linear utilities for agents.

Finally, our approach assumes complete knowledge of all agents of the distributions in the network. The unreasonableness of this full-knowledge assumption is a problem in Bayesian mechanism design more generally, one approach to tackle it being prior-free

[4]On the other hand, the (computationally intractable) approach based on a reduction to MAID equilibrium computation presented in section 4 could also take into account information channels

mechanism design [Har13]. By modelling sparse dependence structure in our model, we tried to make our full-knowledge assumption more reasonable by allowing the model to represent complex arriving information.

Despite these shortcomings, we mention a few avenues for further work:

A first important direction is proving the uniqueness conjecture in section 3. A proof of this result would allow for the unified proof of several results in the literature. In a second step, one could relax the assumptions on agent communication, in particular the notion of *centrality*.

A second direction pertains to approximation algorithms. We reduced a newly introduced computational problem, MAID Equilibrium, to two problems that are themselves hard (one MAID and one large LP). Given that even IDs are only fixed-parameter tractable, as [MCZ12] report, there is not much hope for efficient algorithms. It would be interesting, however, to study approximation algorithms.

The non-existence of non-cooperative MAID equilibria opens up two directions of further research: The first is to define equilibrium concepts that admit an equilibrium and to characterise their graphical properties such as strategic relevance of action nodes (as we did for one commitment agent and the equilibrium concept of MAIDs). The second is to consider the problem of equilibrium existence as a *forbidden subgraph* problem: Is there a set of subgraphs whose containment in MAIDs is equivalent to equilibrium non-existence for some parametrisation?

Moreover, our LP reduction allows for the generalisation of some results from the economics literature that do not use functional specifications so thoroughly that we would not expect generalisations in more general models. The first result is [BS01], which characterises *partial commitment* (which we did not permit in section 5). A generalisation of this result would consist in proving that the strategy characterisation we proved in section 5 also holds with non-zero probability for choice nodes that are not commitment nodes. Similarly, exploiting the structure of the LP one could generalise [CM88] by showing that some of the inequalities in the LP are redundant when agents report the same or correlated quantities.

Finally, mechanism design is lacking concise graphical representations. Indeed, papers rarely show illustrations, and if so, then precise semantics of them are lacking. We hope that MAIDs *with commitment* will allow for illustrations with precise semantics and help students and researchers to more easily grasp mechanism design theory.

References

[BEW] Nils Bertschinger, Olbrich Eckehard, and David H Wolpert. „When more information processors in a distributed system can(not) help“.

[Bis06] Christopher M Bishop. *Pattern recognition and machine learning*. Springer, 2006.

[Bla53] David Blackwell. „Equivalent comparisons of experiments“. In: *The annals of mathematical statistics* (1953), pp. 265–272.

[BNS07] Liad Blumrosen, Noam Nisan, and Ilya Segal. „Auctions with severely bounded communication“. In: *Journal of Artificial Intelligence Research* 28 (2007), pp. 233–266.

[Bru] Boudewijn de Bruin. „Game Transformations and Game Equivalence“. Manuscript.

[BS01] Helmut Bester and Roland Strausz. „Contracting with imperfect commitment and the revelation principle: the single agent case“. In: *Econometrica* 69.4 (2001), pp. 1077–1098.

[BV19] Dirk Bergemann and Juuso Välimäki. „Dynamic mechanism design: An introduction“. In: *Journal of Economic Literature* 57.2 (2019), pp. 235–74.

[CM88] Jacques Cremer and Richard P McLean. „Full extraction of the surplus in Bayesian and dominant strategy auctions“. In: *Econometrica: Journal of the Econometric Society* (1988), pp. 1247–1257.

[CS02] Vincent Conitzer and Tuomas Sandholm. „Complexity results about Nash equilibria“. In: *arXiv preprint cs/0205074* (2002).

[CS04] Vincent Conitzer and Tuomas Sandholm. „Computational Criticisms of the Revelation Principle“. In: *Proceedings of the 5th ACM Conference on Electronic Commerce*. 2004, pp. 262–263.

[Eve+19] Tom Everitt et al. „Understanding Agent Incentives using Causal influence diagrams, Part I: Single Action Settings“. In: *arXiv preprint arXiv:1902.09980* (2019).

[GLS93] Martin Grötschel, László Lovász, and Alexander Schrijver. „The Ellipsoid Method“. In: *Geometric Algorithms and Combinatorial Optimization*. Springer, 1993, pp. 64–101.

[Har13] Jason D Hartline. „Mechanism design and approximation“. Book Draft. 2013.

[How66] Ronald A Howard. „Information value theory“. In: *IEEE Transactions on systems science and cybernetics* 2.1 (1966), pp. 22–26.

[Jac10] Matthew O Jackson. *Social and economic networks*. Princeton university press, 2010.

[JLB11] Albert Xin Jiang, Kevin Leyton-Brown, and Navin AR Bhat. „Action-graph games“. In: *Games and Economic Behavior* 71.1 (2011), pp. 141–173.

[KFB09] Daphne Koller, Nir Friedman, and Francis Bach. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.

[KM03] Daphne Koller and Brian Milch. „Multi-agent influence diagrams for representing and solving games“. In: *Games and economic behavior* 45.1 (2003), pp. 181–221.

[KM86] Elon Kohlberg and Jean-Francois Mertens. „On the strategic stability of equilibria“. In: *Econometrica: Journal of the Econometric Society* (1986), pp. 1003–1037.

[LF87] Irving H Lavalle and Peter C Fishburn. „Equivalent decision trees and their associated strategy sets“. In: *Theory and Decision* 23.1 (1987), pp. 37–63.

[LN01] Steffen L Lauritzen and Dennis Nilsson. „Representing and solving decision problems with limited information“. In: *Management Science* 47.9 (2001), pp. 1235–1251.

[MCZ12] Denis Deratani Mauá, Cassio Polpo de Campos, and Marco Zaffalon. „Solving limited memory influence diagrams“. In: *Journal of Artificial Intelligence Research* 44 (2012), pp. 97–140.

[Mil+08] Brian Milch et al. *Ignorable information in multiagent scenarios*. Tech. rep. 2008.

[MM96] Richard D McKelvey and Andrew McLennan. „Computation of equilibria in finite games". In: *Handbook of computational economics* 1 (1996), pp. 87–142.

[MWG95] Andreu Mas-Colell, Michael Dennis Whinston, and Jerry R Green. *Microeconomic theory*. New York: Oxford university press, 1995.

[Mye81] Roger B Myerson. „Optimal auction design". In: *Mathematics of operations research* 6.1 (1981), pp. 58–73.

[Mye82] Roger B Myerson. „Optimal coordination mechanisms in generalized principal–agent problems". In: *Journal of mathematical economics* 10.1 (1982), pp. 67–81.

[Mye86] Roger B Myerson. „Multistage games with communication". In: *Econometrica: Journal of the Econometric Society* (1986), pp. 323–358.

[Nis+07] Noam Nisan et al. *Algorithmic game theory*. Cambridge university press, 2007.

[Nis15] Noam Nisan. „Algorithmic mechanism design: through the lens of multiunit auctions". In: *Handbook of Game Theory with Economic Applications*. Vol. 4. Elsevier, 2015, pp. 477–515.

[NS06] Noam Nisan and Ilya Segal. „The communication requirements of efficient allocations and supporting prices". In: *Journal of Economic Theory* 129.1 (2006), pp. 192–224.

[PD96] Judea Pearl and Rina Dechter. „Identifying Independencies in causal graphs with Feedback". In: *Proceedings of the Twelfth International Conference on Uncertainty in Artificial Intelligence*. 1996, pp. 420–426.

[SH10] Ross Shachter and David Heckerman. *Pearl causality and the value of control*. 2010.

[SW17] Takuo Sugaya and Alexander Wolitzky. *Revelation Principles in Multistage Games*. Tech. rep. Stanford University and MIT, 2017.

[Tho52] F. Thompson. *Equivalence of games in extensive form*. Tech. rep. RM 759. The Rand Corporation, 1952.

[Yam10] Takuro Yamashita. „Mechanism games with multiple principals and three or more agents". In: *Econometrica* 78.2 (2010), pp. 791–801.