Shapley Revisited: Tractable Responsibility Measures for Query Answers

MEGHYN BIENVENU, Univ. Bordeaux, CNRS, Bordeaux INP, LaBRI, UMR 5800, France
DIEGO FIGUEIRA, Univ. Bordeaux, CNRS, Bordeaux INP, LaBRI, UMR 5800, France
PIERRE LAFOURCADE, Univ. Bordeaux, CNRS, Bordeaux INP, LaBRI, UMR 5800, France

The Shapley value, originating from cooperative game theory, has been employed to define responsibility measures that quantify the contributions of database facts to obtaining a given query answer. For non-numeric queries, this is done by considering a cooperative game whose players are the facts and whose wealth function assigns 1 or 0 to each subset of the database, depending on whether the query answer holds in the given subset. While conceptually simple, this approach suffers from a notable drawback: the problem of computing such Shapley values is #P-hard in data complexity, even for simple conjunctive queries. This motivates us to revisit the question of what constitutes a reasonable responsibility measure and to introduce a new family of responsibility measures – weighted sums of minimal supports (WSMS) – which satisfy intuitive properties. Interestingly, while the definition of WSMSs is simple and bears no obvious resemblance to the Shapley value formula, we prove that every WSMS measure can be equivalently seen as the Shapley value of a suitably defined cooperative game. Moreover, WSMS measures enjoy tractable data complexity for a large class of queries, including all unions of conjunctive queries. We further explore the combined complexity of WSMS computation and establish (in)tractability results for various subclasses of conjunctive queries.

CCS Concepts: • **Theory of computation → Data provenance**; Incomplete, inconsistent, and uncertain databases; Logic and databases; • **Information systems →** *Relational database model*.

Additional Key Words and Phrases: Shapley value, numeric responsibility measures, minimal supports, non-numeric queries, monotone queries, conjunctive queries

ACM Reference Format:
Meghyn Bienvenu, Diego Figueira, and Pierre Lafourcade. 2025. Shapley Revisited: Tractable Responsibility Measures for Query Answers. *Proc. ACM Manag. Data* 3, 2 (PODS), Article 112 (May 2025), 26 pages. https://doi.org/10.1145/3725249

☞ This pdf contains internal links: clicking on a notion leads to its *definition*.
☞ A long version of this paper with detailed proofs is available at https://arxiv.org/abs/2503.22358.

1 Introduction

Responsibility measures assign scores to database facts based upon how much they contribute to the obtention of a given query answer, thereby providing a quantitative notion of explanation for query results. In the present paper, we shall take a fresh look at how to define and compute such responsibility measures, focusing on the class of monotone non-numeric queries, which notably includes (unions of) conjunctive queries and path queries. Although several responsibility measures

Authors' Contact Information: Meghyn Bienvenu, meghyn.bienvenu@cnrs.fr, Univ. Bordeaux, CNRS, Bordeaux INP, LaBRI, UMR 5800, F-33400, Talence, France; Diego Figueira, diego.figueira@cnrs.fr, Univ. Bordeaux, CNRS, Bordeaux INP, LaBRI, UMR 5800, F-33400, Talence, France; Pierre Lafourcade, pierre.lafourcade@u-bordeaux.fr, Univ. Bordeaux, CNRS, Bordeaux INP, LaBRI, UMR 5800, F-33400, Talence, France.

ACM 2836-6573/2025/5-ART112
https://doi.org/10.1145/3725249

have been explored in the database literature [1, 17, 21, 24], one particular measure, based upon the Shapley value, has been the focus of much of the recent work [2, 4, 13–15, 17, 23].

We recall that the Shapley value [25] was first introduced in game theory for the purpose of fairly distributing wealth amongst the players in a cooperative game. Formally, such a game is defined using a "wealth function" $\omega : 2^{\text{Players}} \rightarrow \mathbb{R}_{\geqslant 0}$ which assigns a number to every subset of players. The Shapley value then defines a function ϕ_ω : Players $\rightarrow \mathbb{R}$ from players to numbers, which measures their individual contributions. The Shapley function has been famously characterized as the only function satisfying four fundamental axioms (which will be detailed in Section 3).

In the database context, the Shapley value has been utilized for distributing the responsibility for a given query answer among database facts. This is done by considering games in which the facts are the players and the query is captured using the wealth function. For *numeric queries* (*i.e.*, queries yielding a number when evaluated on a database), the sensible choice for the wealth function is the query itself [17]. That is, for a given input numeric query q, database $\mathcal{D}$, and set of facts $S \subseteq \mathcal{D}$ over which we want to distribute the contribution – known as 'endogenous' facts for historical reasons –, the Shapley value distributes the total wealth $q(\mathcal{D}) = N$ among the facts of S (*i.e.*, the 'players') using q itself as the wealth function.

For non-numeric queries, whose output is a set of tuples of constants rather than a number, it is less obvious how to define the wealth function. The solution that was originally proposed in [17] and adopted in all subsequent studies is to first reduce to the Boolean case by considering the Boolean query $q(\bar{a})$ associated with the input query q and answer tuple $\bar{a}$, and then to treat Boolean queries as numeric queries that output 0 or 1 depending on whether the query holds true in a set of facts. We shall refer to the Shapley value obtaining by encoding queries using such a 0/1 wealth function as the *drastic* Shapley value.

The use of the drastic Shapley value for responsibility attribution of non-numeric queries has received significant attention of late [2, 4, 13–15, 17, 23]. The main focus has been on precisely identifying the data complexity of computing the drastic Shapley value, considering various classes of queries and drawing connections to the related tasks of probabilistic query evaluation and model counting. The key takeaway is that drastic Shapley value computation is computationally challenging, being $\mathsf{FP}^{\#\mathsf{P}}$-complete in data complexity even for some very simple conjunctive queries. The situation is more favourable if one moves from exact to approximate computation, namely, Monte-Carlo sampling provides a fully polynomial randomized approximation scheme (FPRAS) for every UCQ [17, Corollary 4.13], even if practical implementations appear to be challenging (see experiments and discussions in [4, §6.2]).

Another potential drawback to the drastic Shapley value is that it may be hard to interpret the resulting scores, in particular to understand how a change in the database would impact the relative scores of a pair of facts, which may be influenced by seemingly unrelated facts (as exemplified by Example 4.1). In its original conception, the Shapley value uses a wealth function whose numbers correspond to some 'cost' or 'wealth' in some implicit *unit of measure* (*e.g.* dollars, number of workers, number of tuples satisfying some property). This same idea carries over to numeric queries, as the 'total wealth' to distribute is given by the query itself, which outputs a number with a clear meaning. However, the drastic Shapley approach casts the Boolean values 'true' and 'false' into the (real) numbers 1 and 0, in order to apply operations like mean. This can make the interpretation of the drastic Shapley value not always intuitive to a non-expert end user, which might hence hinder its use for explainability tasks.

Given both the high computational complexity and non-obvious interpretability of the drastic Shapley value, a natural question is whether we can find another responsibility measure for non-numeric queries that is computationally and conceptually simpler. One might be tempted to reply 'no', arguing that the uniqueness result for the Shapley axioms forces us to adopt the drastic Shapley value. However, these axioms were introduced for cooperative games, rather than responsibility measures for queries, and their meaning crucially depends on the manner in which queries are translated into games. In fact, we shall see that when these axioms are translated from games to responsibility measures, using the drastic wealth function, only two of the four axioms yield unarguably desirable properties for such measures (and one axiom is actually nonsensical).

In view of the preceding discussion, it is both natural and possible to explore alternative responsibility measures for non-numeric queries. This will be the main topic of the present paper, whose conceptual and technical contributions are summarized in what follows.

Our first conceptual contribution is to propose a set of *desirable properties for responsibility measures*, focusing on the case of monotone non-numeric queries. We begin by examining in Section 3 the properties that result from mapping the Shapley axioms into the database setting using the drastic Shapley wealth function. We observe that only two yield clearly desirable properties in our setting, while the other two are either nonsensical or of debatable interest. As this leaves us with only two, rather weak, properties, we then introduce in Section 4.1 two additional properties that capture intuitions as to how a reasonable responsibility measure should behave.

Our second conceptual contribution is to propose in Section 4.2 a *novel family of responsibility measures* specifically tailored to monotone non-numeric queries, which we name *weighted sums of minimal supports* (*WSMS*). As the name suggests, the measure of a fact consists of a weighted sum of the sizes of minimal supports of the query (answer) that contain the fact. These measures have an intuitive definition, they enjoy the identified desired properties, and the values arising from these measures have a simple interpretation. Further, we can show that each measure from this family can be equivalently seen as the Shapley value of a suitable cooperative game (Section 4.3). We also argue that the family is flexible and can accommodate weight functions giving more importance to small supports than to the number of supports, or vice-versa (Section 4.4).

Our main technical contribution is a *complexity analysis of WSMS computation*. In Section 5, we show that, in contrast to the drastic Shapley value, WSMSs enjoy tractable data complexity for all unions of conjunctive queries (and more generally, bounded queries) via a straightforward algorithm. For combined complexity, we obtain a general $\mathsf{FP}^{\#\mathsf{P}}$ bound for a wide range of query classes (Section 6.1). We then focus on conjunctive queries (CQs) and show that, unfortunately, even for acyclic CQs without constants, computing WSMS is #P-hard. However, for acyclic self-join free CQs, the WSMS computation can be reduced to evaluating 'counting CQs' (more precisely, counting the number of homomorphisms from acyclic CQs), which is known to be tractable. Further, in Section 6.3 we show that if both the generalized hypertree width and the amount of self-joins (via a measure we define with this purpose) are bounded, then WSMSs on CQs can be computed in polynomial time.

As a final contribution, we explore the *properties of some related measures*. First, having observed that WSMSs are closely connected to the MI inconsistency measure [12], we define in Section 7.1 three new responsibility measures based upon inconsistency measures and show that they are less suitable for semantic or computational reasons. We likewise show in Section 7.2 that redefining WSMSs by counting homomorphisms rather than minimal supports yields a measure which violates a basic property. Finally, in Section 7.3, we recall the SHAP score that has been used for explaining classifiers in machine learning and suggest how the WSMS definition could be adapted to fix some semantic problems with the SHAP score (albeit at the cost of a higher computational complexity).

2 Preliminaries

We fix disjoint infinite sets Const, Var of *constants* and *variables*, respectively. For any syntactic object O (*e.g.* database, query), we will use $var(O)$ and $const(O)$ to denote the sets of variables and constants contained in O, and let $term(O) \hat{=} var(O) \cup const(O)$ denote its set of *terms*.

A *(relational) schema* is a finite set of relation symbols, each associated with a (positive) arity. A *(relational) atom* over a schema Σ takes the form $R(\bar{t})$ where R is a *relation name* from Σ of some arity k, and $\bar{t} \in (\text{Const} \cup \text{Var})^k$. A *fact* is an atom which contains only constants. A *database* $\mathcal{D}$ over a schema Σ is a finite set of facts over Σ.

A *(non-numeric) query* q of arity $k \geqslant 0$ over schema Σ is a syntactic object associated with a *semantics* consisting of a function that maps every database $\mathcal{D}$ over Σ to a set $q(\mathcal{D})$ of k-tuples of elements from $const(\mathcal{D})$,[1] called the *answers* to q on $\mathcal{D}$. When $k = 0$, we say that q is a *Boolean query*, and we shall write $\mathcal{D} \models q$ in place of $\mathcal{D}(q) = \{()\}$. As explained in Section 3.1, it will be sufficient to work with Boolean queries, so *except where otherwise indicated, we will henceforth abuse terminology and write query to mean Boolean query*.

We shall further focus our attention on *monotone* queries, i.e. queries q such that $\mathcal{D} \models q$ implies $\mathcal{D} \cup \mathcal{D}' \models q$ for every pair of databases $\mathcal{D}, \mathcal{D}'$. When $\mathcal{D} \models q$ and q is monotone, we call $\mathcal{D}$ a *support* for q and say it is a *minimal support* if $\mathcal{D}$ properly contains no other support. We denote by $\mathsf{MS}_q(\mathcal{D})$ the set of all minimal supports for q that are subsets of $\mathcal{D}$. A monotone query q is *bounded* if the sizes of minimal supports for q are bounded by a constant (independent of the database). We say that a fact $\alpha \in \mathcal{D}$ is *relevant* for q in $\mathcal{D}$ if $\alpha \in S$ for some $S \in \mathsf{MS}_q(\mathcal{D})$, and *irrelevant* otherwise.

We now define some concrete classes of monotone (Boolean) queries. A *conjunctive query* (*CQ*) over Σ takes the form $q = \exists \bar{x}\, \varphi$ where φ is a conjunction of atoms over Σ and $\bar{x} = var(\varphi)$. We use $atoms(q)$ to denote the set of atoms of a CQ q. An *acyclic CQ* (ACQ) is a CQ $q = \exists \bar{x}\, \varphi$ such that there is a tree whose vertices are $atoms(q)$ and such that for vertices α, α', β, $var(\alpha) \cap var(\alpha') \subseteq var(\beta)$ whenever β lies on the unique simple path between α, α'. We shall denote by CQ (resp. ACQ) the class of CQs (resp. ACQs). Other prominent classes of monotone queries include: *unions of conjunctive queries* (*UCQs*), which are finite disjunctions of CQs, and *path queries*, which take the form $\mathcal{L}(c, d)$ where $\mathcal{L}$ is a language over a set of binary relation names and c, d are constants, with the semantics of testing whether there is a directed path from c to d whose label is in $\mathcal{L}$ in the input database. In particular, a (Boolean) *regular path query* (*RPQ*) is a path query where $\mathcal{L}$ is a regular language.

We recall that the semantics of a CQ q is defined by $\mathcal{D} \models q$ iff there is a function $h : term(q) \longrightarrow term(\mathcal{D})$ such that (i) $R(\bar{t}) \in atoms(q)$ implies $R(h(\bar{t})) \in \mathcal{D}$, and (ii) $h(c) = c$ for every $c \in const(q)$. Such a function is called a (query) *homomorphism* (from q to $\mathcal{D}$), and we write $q \xrightarrow{\text{hom}} \mathcal{D}$ to indicate the existence of such a homomorphism. We shall also consider homomorphisms of a CQ q to another CQ q' , denoted $q \xrightarrow{\text{hom}} q'$ and defined similarly. The *core* of a CQ q, denoted $core(q)$, is the unique (up to isomorphism) minimal equivalent query [3].

We use the standard notion of *polynomial-time Turing-reductions* between numeric problems Ψ_1, Ψ_2, and we write $\Psi_1 \leq_{poly} \Psi_2$ if there is a polynomial-time algorithm for Ψ_1 using a Ψ_2 oracle.

3 Responsibility Measures and the Drastic Shapley Value

The (drastic) Shapley value has been widely studied as a way of defining a responsibility measure for database queries. A seminal and oft-cited result [25] in cooperative game theory establishes that the Shapley value is the unique wealth distribution measure that satisfies a specific set of 'Shapley axioms'. It is tempting to think that this unicity result transfers to the setting of database

[1] Note the common notation between the syntax q and the (implicit) semantics $\mathcal{D} \mapsto q(\mathcal{D})$.

responsibility measures, and thus that the drastic Shapley value considered in prior work is the only reasonable approach. However, we consider that there is no single valid notion of a responsibility measure, and we shall show in this section that the Shapley axioms are insufficient to argue the contrary. We start by introducing the notion of responsibility measure for non-numeric database queries, before recalling the definition of the Shapley value from game theory and how it has been utilized to define a concrete such responsibility measure. Finally, we discuss the relevance of the Shapley axioms in the database setting and motivate the interest of exploring alternative measures.

3.1 Responsibility Measures

In line with prior studies [2, 13, 17, 21], we work with *partitioned databases*, in which the database $\mathcal{D}$ is partitioned into *endogenous* and *exogenous* facts, denoted by $\mathcal{D}_{\mathsf{n}}$ and $\mathcal{D}_{\mathsf{x}}$ respectively. The aim will be to quantify the responsibility of a given endogenous fact to the obtention of some query answer. Formally, we shall be interested in *responsibility measures*, defined as functions ϕ which take as input a partitioned database $\mathcal{D}$, a (possibly non-Boolean) query q, an answer $\bar{a}$ to q in $\mathcal{D}$, and a fact $\alpha \in \mathcal{D}_{\mathsf{n}}$, and which output a quantitative score measuring how much α contributes to $\bar{a} \in q(\mathcal{D})$. We may however simplify the presentation by replacing the input answer $\bar{a}$ and (possibly non-Boolean) query q by the associated Boolean query $q(\bar{a})$, defined by letting $\mathcal{D} \models q(\bar{a})$ iff $\bar{a} \in q(\mathcal{D})$. In this manner, we can eliminate the answer tuple from the arguments of the responsibility measure and work instead with ternary responsibility measures, i.e. $\phi(\mathcal{D}, q, \alpha)$ with Boolean q. Henceforth, we shall thus assume w.l.o.g. that the input query is always Boolean, as well as monotone.[2] We discuss in Sections 3.3 and 4.1 the desirable properties of responsibility measures.

3.2 Shapley Value: From Games to Databases

A *cooperative game* is given by a finite set of players P and a *wealth function* $\xi : 2^P \to \mathbb{Q}$ that represents the wealth generated by every possible coalition of players, such that $\xi(\varnothing) = 0$. To define the Shapley value, assume players arrive one by one in a random order, and each earns the difference between the coalition's wealth before she arrives and after. The *Shapley value* of a player $p \in P$ is her expected earnings in this scenario, which can be expressed as:

$$\mathrm{Sh}(P, \xi, p) \triangleq \frac{1}{|P|!} \sum_{\sigma \in \mathbb{P}(P)} \big(\xi(\sigma_{<p} \cup \{p\}) - \xi(\sigma_{<p})\big) \tag{1}$$

where $\mathbb{P}(P)$ denotes the set of permutations of P and $\sigma_{<p}$ the set of players that appear strictly before p in the permutation σ.

Besides the Shapley value, another wealth distribution scheme, the *Banzhaf power index*, has been applied to responsibility distribution for database queries [1, 24].[3] This alternative scheme is closely related to the Shapley value, to the point where [14] introduced the notion of *Shapley-like* scores to collectively designate a class of wealth distribution schemes that contains them both. Formally, Shapley-like scores are defined by the formula:

$$\mathcal{S}_c(P, \xi, p) \triangleq \sum_{B \subseteq P \setminus \{p\}} c(|B|, |P|)(\xi(B \cup \{p\}) - \xi(B)) \tag{2}$$

where $c(|B|, |P|)$ is a *coefficient function* characterising the measure. The Shapley value is obtained by letting $c(|B|, |P|) = \frac{|B|!(|P|-|B|-1)!}{|P|!}$ and the Banzhaf power index by letting $c(|B|, |P|) = 1$.

[2]The study of non-monotone queries is left for future work, see Section 8.

[3]The "causal effect" measure considered in [24] coincides with the Banzhaf power index [17, Proposition 5.5].

To employ these notions in the database setting, one must specify how to build games from queries and databases. Concretely, we need a *wealth function family* $\Xi^\star$ which takes as parameters a query q and a set $\mathcal{D}_\mathsf{x}$ of exogenous facts and defines a wealth function $\xi^\star_{q,\mathcal{D}_\mathsf{x}}$. For instance the standard wealth function family in the literature [2, 17] is what we call the *drastic* wealth function family,[4] denoted by Ξ^{dr} and defined as follows: $\xi^{\mathsf{dr}}_{q,\mathcal{D}_\mathsf{x}}$ assigns 1 to every set S of facts such that $S \cup \mathcal{D}_\mathsf{x} \models q$ and 0 to the others, unless $\mathcal{D}_\mathsf{x} \models q$ in which case it assigns 0 to all sets.

Once a wealth function family has been defined, it can be combined with the Shapley value to obtain a responsibility measure, whose output is $\mathrm{Sh}(\mathcal{D}_\mathsf{n}, \xi^\star_{q,\mathcal{D}_\mathsf{x}}, \alpha)$. Given a class of queries $\mathcal{C}$, we denote by $\mathrm{SVC}^\star_\mathcal{C}$ the associated computational problem of computing, for a given query $q \in \mathcal{C}$, partitioned database $\mathcal{D}$ and fact $\alpha \in \mathcal{D}_\mathsf{n}$, the value $\mathrm{Sh}(\mathcal{D}_\mathsf{n}, \xi^\star_{q,\mathcal{D}_\mathsf{x}}, \alpha)$. When $\mathcal{C}$ consists of a single query q, we simply write $\mathrm{SVC}^\star_q$. Note the superscript $\star$ that is consistent across all notations. Prior work on Shapley value computation for Boolean queries considered $\star = \mathsf{dr}$, i.e. $\mathrm{SVC}^{\mathsf{dr}}_q$ and $\mathrm{SVC}^{\mathsf{dr}}_\mathcal{C}$.

3.3 Shapley Axioms

The importance of the Shapley value stems from satisfying several axioms, deemed desirable in many scenarios for fair distribution of wealth or responsibility in cooperative games. The four fundamental axioms put forth by Shapley for a real-valued wealth distribution function ψ (whose inputs are a set of players, a wealth function, and a distinguished player) are as follows:

- (Sym) *Symmetry:* if two players $\alpha, \beta \in P$ are equivalent in the sense that $\forall S \subseteq P \setminus \{\alpha, \beta\}.\ \xi(S \cup \{\alpha\}) = \xi(S \cup \{\beta\})$, then $\psi(P, \xi, \alpha) = \psi(P, \xi, \beta)$;
- (Null) *Null player:* any (so-called *null*) player who does not contribute to increasing the wealth (*i.e.*, such that $\xi(S \cup \{\alpha\}) = \xi(S)$ for all S) must obtain 0 as contribution;
- (Lin) *Linearity:* the value of the sum of two wealth functions $\xi_1 + \xi_2$ is just the sum of values over the separate wealth functions: $\psi(P, \xi_1 + \xi_2, \alpha) = \psi(P, \xi_1, \alpha) + \psi(P, \xi_2, \alpha)$;
- (Eff) *Efficiency:* the sum $\sum_{\alpha \in P} \psi(P, \xi, \alpha)$ of all contributions equals the total wealth $\xi(P)$.

Shapley has famously shown in [25, p. 295] that Sh is the *unique* function ψ satisfying (Sym)–(Eff) and hence that the Shapley value is 'the' function to consider in any context where the axioms above make sense and should be met.[5]

However, the preceding axioms are phrased in terms of cooperative games, whereas we are interested in properties of responsibility measures, which take queries and databases as arguments. The interpretation of the Shapley axioms in the database setting thus crucially depends on the way we model the queries as cooperative games, *i.e.* the wealth function family we use. Given that all prior work has focused on the drastic wealth function, Ξ^{dr}, it is instructive to instantiate the above axioms with Ξ^{dr} in order to understand which properties of responsibility measures have been invoked to justify the drastic Shapley value for non-numeric queries. We first consider (Sym):

(Sym-dr) If two facts $\alpha, \beta \in \mathcal{D}_\mathsf{n}$ are equivalent in the sense that for all $S \subseteq \mathcal{D}_\mathsf{n} \setminus \{\alpha, \beta\}$, $S \cup \{\alpha\} \cup \mathcal{D}_\mathsf{x} \models q$ iff $S \cup \{\beta\} \cup \mathcal{D}_\mathsf{x} \models q$, then $\phi(\mathcal{D}, q, \alpha) = \phi(\mathcal{D}, q, \beta)$.

We consider (Sym-dr) to be a very desirable property. In fact, we would suggest to strengthen it as follows to ensure syntax independence not only for the input facts but also for the queries:

[4]The name *drastic* is borrowed from the closely related *drastic inconsistency measure* (see, *e.g.*, measure 'I_d' in [18, 26]) assigning 0 or 1 depending on whether the database is consistent *w.r.t.* a property. See also Section 7.1.

[5]In [25], Shapley presents only 3 axioms: (Null) is omitted and a different stronger alternative axiom for (Eff) is used, which implies both (Null) and (Eff) as stated here. We prefer to adopt the widely used modern presentation of the Shapley axioms, see for example [16, §2].

(Sym-db) If two queries q_1, q_2 have the same semantics (*i.e.*, define the same Boolean function), and facts $\alpha, \beta \in \mathcal{D}_\mathsf{n}$ are such that for all $S \subseteq \mathcal{D}_\mathsf{n} \setminus \{\alpha, \beta\}$, $S \cup \{\alpha\} \cup \mathcal{D}_\mathsf{x} \models q_1$ iff $S \cup \{\beta\} \cup \mathcal{D}_\mathsf{x} \models q_2$, then $\phi(\mathcal{D}, q_1, \alpha) = \phi(\mathcal{D}, q_2, \beta)$.

Turning now to (Null), we observe that null players translate into irrelevant facts, yielding:

(Null-dr) If a fact $\alpha \in \mathcal{D}_\mathsf{n}$ is irrelevant, then $\phi(\mathcal{D}, q, \alpha) = 0$.

Again, we find (Null-dr) desirable but insufficient, as we wish to be able to decide relevance:

(Null-db) If a fact $\alpha \in \mathcal{D}_\mathsf{n}$ is irrelevant then $\phi(\mathcal{D}, q, \alpha) = 0$; otherwise $\phi(\mathcal{D}, q, \alpha) > 0$.[6]

Moving on to (Lin), we come to our first issue. This axiom should intuitively translate as "the sum of the responsibility for two queries is the responsibility for the sum query", which makes sense when dealing with numeric queries (for which the sum query is well defined), but whose meaning is unclear for non-numeric queries. When instantiated with Ξ^{dr}, the property is nonsensical because for all non-trivial cases $\xi^{\mathsf{dr}}_{q_1,\mathcal{D}_\mathsf{x}}(\mathcal{D}_\mathsf{n}) + \xi^{\mathsf{dr}}_{q_2,\mathcal{D}_\mathsf{x}}(\mathcal{D}_\mathsf{n}) = 2$, which cannot coincide with $\xi^{\mathsf{dr}}_{q,\mathcal{D}_\mathsf{x}}(\mathcal{D}_\mathsf{n})$ (whose value is 0 or 1), no matter which Boolean q we use to represent the 'sum query'.[7]

Let us finally consider the meaning of (Eff) once instantiated with the drastic wealth function:

(Eff-dr) The sum $\sum_{\alpha \in \mathcal{D}_\mathsf{n}} \phi(\mathcal{D}, q, \alpha) = 1$ if $\mathcal{D} \models q$ and $\mathcal{D}_\mathsf{x} \not\models q$, and 0 otherwise.

While it seems reasonable enough to ask that the sum of contributions of the facts is equal to the value attained by the full set of endogenous facts, it is rather more questionable why we should always arrive at the same total value of 1 when $\mathcal{D} \models q$ and $\mathcal{D}_\mathsf{x} \not\models q$, irrespective of the particular q and $\mathcal{D}$. The fact that queries are modelled using constant-sum games means that the drastic Shapely value is highly sensitive to database changes: the addition of *any* new relevant fact will affect the values of *all* relevant facts, as they need to share the same total wealth (of 1). This, in turn, can make it challenging to interpret the drastic Shapley value of facts without appealing to the probabilistic process that underlies Equation (1). Note that this differs markedly from the case of numeric queries for which the total value – defined via the query result – varies according to the input database and finds its meaning in the query itself.

In view of the identified issues with the translations of (Lin) and (Eff) into the database setting, the Shapley axioms cannot be used to argue that the drastic Shapley value is the only (or the best) approach to defining a reasonable responsibility measure. Indeed, only (Sym-db) and (Null-db) are truly natural, and they constitute rather weak minimal requirements.

This raises the obvious questions: *What are desirable properties for responsibility measures of non-numeric queries? And how can they be realized?* In Section 4, we address these questions and propose a novel family of measures conforming to the identified properties.

4 Alternative responsibility measures

As discussed in Section 3, the Shapley axioms provide insufficient justification for arguing that the drastic Shapley value is the *only reasonable* or the *best* measure for non-numeric queries.Further, $\mathsf{SVC}^{\mathsf{dr}}_q$ is #P-hard for some very simple CQs [17]. This motivates us to consider alternative measures, which can be done in two ways:

Approach 1: we may consider alternative responsibility attribution functions more specifically tailored to monotone non-numeric queries which may not be based on a Shapley score, or

Approach 2: we may insist in the use of the Shapley value but for a different associated wealth function, one in which axioms (Eff) and (Lin) are more meaningful.

[6] Recall that the query q we consider is always Boolean and monotone.

[7] An alternative to the linearity axiom that still makes sense for Boolean games has been introduced in [5, Axiom S3'], but we believe it still doesn't yield a very natural axiom for responsibility measures. See Remark A.1.

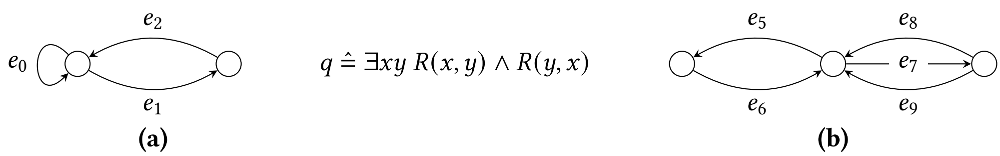

Fig. 1. Pathological examples for some responsibility measures. The databases are represented as graphs with nodes being constants and edges being R-tuples.

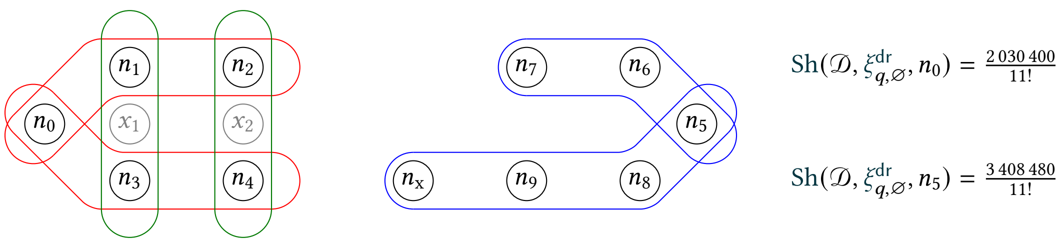

Fig. 2. Instance for Example 4.1. $\mathcal{D}_{\mathrm{n}} \mathrel{\hat{=}} \{n_1, \ldots, n_{\mathrm{x}}\}$, $\mathcal{D}_{\mathrm{x}} \mathrel{\hat{=}} \{x_1, x_2\}$, and the query q is such that the minimal supports are the colored regions (the red ones contain e_0, the blue ones e_5 and the green ones neither).

This section is organized as follows. In Section 4.1 we follow the first approach by identifying a set of truly desirable properties and proposing an alternative family of measures conforming to them in Section 4.2. Then in Section 4.3 we show that these alternative measures can all be seen as Shapley values for different wealth functions, and thus that they could also have been found by following the second approach. Finally Section 4.4 will suggest explicit measures within the family.

4.1 Desired Properties of Responsibility Measures

We already have from Section 3 the properties (Sym-db) and (Null-db) which capture some basic requirements. To identify further desirable properties, we shall go back to [17] where the drastic Shapley value was first introduced. There, a previously defined quantitative measure (the causal responsibility [21]) was argued to yield counterintuitive results on a specific example [17, Example 5.2], a slight adaptation of which is presented in Figure 1-(a). Indeed, the causal responsibility assigns the same score to all edges in this example, even though e_0 intuitively has more responsibility as it satisfies the query on its own, whereas e_1 and e_2 must be used together. More generally, this idea can be captured by the following axiom (*AOTBE* stands for "All Other Things Being Equal"):

(MS1) AOTBE, a fact that appears in smaller minimal supports should have higher responsibility.

In a similar vein, it turns out that the causal responsibility also gives the same score to e_5 and e_7 in the example represented in Figure 1-(b). Again, this is counterintuitive because the fact e_5 only appears in a single minimal support ($\{e_5, e_6\}$) while e_7 appears in two ($\{e_7, e_8\}$ and $\{e_7, e_9\}$). This suggests the following more general property:

(MS2) AOTBE, a fact that appears in more minimal supports should have higher responsibility.

The axioms (MS1) and (MS2) have been stated informally. While mathematically precise formulations can be achieved and can be found in Appendix B.1, they turn out to be quite technical as the notion of "AOTBE" is non-trivial to formalize. For instance, the most naïve definitions will fail to predict the behavior of the drastic Shapley value on the following example.

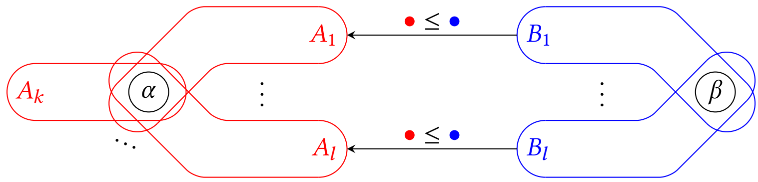

Fig. 3. A visual representation of the setting considered by (MStest). The A_i, B_j are the minimal supports of q.

Example 4.1. Consider the instance depicted in Figure 2. Fact n_0 appears in two minimal supports of size 3 (in red) while n_5 appears in one of size 3 and one of size 4 (in blue). However, the situation isn't entirely symmetrical between n_0 and n_5 because the minimal supports that contain neither of them (in green) interfere with n_0 only. As it turns out, the drastic Shapley value will consider that this difference is sufficient to attribute a smaller value to n_0. Note that if one were to remove x_1 and x_2, the green minimal supports would disappear and n_0 would get the highest value. This may be unexpected to a non-expert end user, since these two facts don't share any minimal support with either n_0 nor n_5. △

In order to obtain a concrete criterion that is not too complex, we shall focus on a particular class of scenarios – depicted in Figure 3 – in which there are no non-trivial interactions between the minimal supports. The fact that any two minimal supports may only intersect on $\{\alpha\}$ or $\{\beta\}$ prevents such unwanted interactions, while the condition that $|A_i| \leq |B_i|$ for all i ensures a loose domination of α over β. By requiring that this domination is strict in some sense, we finally obtain the following necessary condition for a responsibility measure ϕ. Note that although this condition may seem very specific, it already suffices to cover the cases in Figure 1.

(MStest) Let $A_1, \ldots, A_k, B_1, \ldots, B_l$ be sets of facts, with $k \geq l$, such that for all i, j $A_i \cap B_j = \emptyset$, for all i, j with $i \neq j$ $A_i \cap A_j = \{\alpha\}$, $B_i \cap B_j = \{\beta\}$ and $|A_i| \leq |B_i|$. Consider an instance on the input database $\mathcal{D} = A_1 \cup \cdots \cup A_k \cup B_1 \cup \cdots \cup B_l$ where q is such that the minimal supports are the A_i and B_j. If furthermore $k > l$ or $\exists i. |A_i| < |B_i|$, then $\phi(\mathcal{D}, q, \alpha) > \phi(\mathcal{D}, q, \beta)$.

4.2 Weighted Sums of Minimal Supports

Given that minimal supports are central to (MS1), (MS2) and also (Null-db), it seems natural to employ this notion directly in the definition of responsibility measures. We shall thus introduce the *weighted sum of minimal supports* (or *WSMS*), as the quantitative responsibility measure

$$\phi^{w}_{\mathrm{wsms}}(\mathcal{D}, q, \alpha) \doteq \sum_{\substack{S \in \mathrm{MS}_q(\mathcal{D}) \\ \alpha \in S}} w(|S|, |\mathcal{D}|) \quad \text{if } \mathcal{D}_{\mathrm{x}} \not\models q, \text{ or } 0 \text{ otherwise} \tag{3}$$

for some *weight function* $w : \mathbb{N} \times \mathbb{N} \to \mathbb{R}$. We say that w is a *tractable weight function* if it can further be computed in polynomial time. Note the special case for $\mathcal{D}_{\mathrm{x}} \models q$, like the one for Ξ^{dr}.

Proposition 4.2. *If w is positive and strictly decreasing (i.e., such that $k < l \Rightarrow w(k, n) > w(l, n)$ for every n), then ϕ^{w}_{wsms} satisfies (Sym-db), (Null-db) and (MStest). Further, any drastic Shapley-like score with positive coefficient function satisfies these axioms as well.*[8]

Arguably the most natural positive and strictly decreasing function is the *inverse weight function* invw : $k, s \mapsto 1/k$. As Proposition 4.3 shows, $\phi^{\mathrm{invw}}_{\mathrm{wsms}}$ coincides with the Shapley value of the wealth function family Ξ^{ms}, defined by $\xi^{\mathrm{ms}}_{q,\mathcal{D}_{\mathrm{x}}}(\mathcal{D}_{\mathrm{n}}) = \sum_{S \in \mathrm{MS}_q(\mathcal{D}_{\mathrm{n}} \uplus \mathcal{D}_{\mathrm{x}})} \frac{|S \cap \mathcal{D}_{\mathrm{n}}|}{|S|}$ if $\mathcal{D}_{\mathrm{x}} \not\models q$, and 0 otherwise. In

[8] All these measures also satisfy (the formal versions of) (MS1) and (MS2) by Propositions B.1 and B.2 in the appendix.

the non-trivial case where $\mathcal{D}_{\mathsf{x}} \models q$, a minimal support $S \subseteq \mathcal{D}_{\mathsf{n}}$ contributes 1 to the total, while a minimal support that is, say, $1/3$ exogenous, only contributes $2/3$. On purely endogenous databases, the formula therefore simplifies to $\xi^{\mathrm{ms}}_{q,\varnothing}(\mathcal{D}_{\mathsf{n}}) = |\mathsf{MS}_q(\mathcal{D}_{\mathsf{n}})|$. Intuitively, this measure considers that the fact that $\mathcal{D} \models q$ is more robust the more minimal supports for q can be found in $\mathcal{D}$. A similar idea underlies the so-called *MI inconsistency measure* [12] (see Section 7 for more details).

Proposition 4.3. *The MS Shapley value, that is, the Shapley value of* Ξ^{ms}*, satisfies the formula:*

$$\mathrm{Sh}(\mathcal{D}_{\mathsf{n}}, \xi^{ms}_{q,\mathcal{D}_{\mathsf{x}}}, \alpha) = \phi^{\mathrm{invw}}_{\mathsf{wsms}}(\mathcal{D}, q, \alpha) = \sum_{\substack{S \in \mathsf{MS}_q(\mathcal{D}) \\ \alpha \in S}} \frac{1}{|S|} \quad \textit{if } \mathcal{D}_{\mathsf{x}} \not\models q, \textit{ or } 0 \textit{ otherwise.}$$

Proof. Consider a query q and a partitioned database $\mathcal{D} \doteq \mathcal{D}_{\mathsf{n}} \uplus \mathcal{D}_{\mathsf{x}}$. For every $S \in \mathsf{MS}_q(\mathcal{D})$ define the query $q_{|S}$ whose only minimal support is S, that is $B \models q_{|S}$ iff $S \subseteq B$. It can be easily seen that $\xi^{\mathrm{ms}}_{q,\mathcal{D}_{\mathsf{x}}}$ on $\mathcal{D}$ can be expressed as $\xi^{\mathrm{ms}}_{q,\mathcal{D}_{\mathsf{x}}} = \sum_{S \in \mathsf{MS}_q(\mathcal{D})} \xi^{\mathrm{ms}}_{q_{|S},\mathcal{D}_{\mathsf{x}}}$ if $\mathcal{D}_{\mathsf{x}} \not\models q$, 0 otherwise.

Since the other case is trivial, assume $\mathcal{D}_{\mathsf{x}} \not\models q$ and consider some $S \in \mathsf{MS}_q(\mathcal{D})$. By (Null), $\mathrm{Sh}(\mathcal{D}_{\mathsf{n}}, \xi^{\mathrm{ms}}_{q_{|S},\mathcal{D}_{\mathsf{x}}}, \alpha) = 0$ for any endogenous fact α outside S. By (Sym), the remaining endogenous facts (*i.e.* the elements of $S \cap \mathcal{D}_{\mathsf{n}}$) must all have the same value, and their sum is fixed to $\frac{|S \cap \mathcal{D}_{\mathsf{n}}|}{|S|}$ by (Eff) (recall we assume $\mathcal{D}_{\mathsf{x}} \not\models q$). Overall this means that $\mathrm{Sh}(\mathcal{D}_{\mathsf{n}}, \xi^{\mathrm{ms}}_{q_{|S},\mathcal{D}_{\mathsf{x}}}, \alpha) = 1/|S|$ if $\alpha \in S$ and 0 otherwise. Note that, so long as $\mathcal{D}_{\mathsf{x}} \not\models q$, this value isn't affected by the distribution of exogenous facts. Finally by (Lin) we can sum the contributions of all subgames to get the desired formula. □

Proposition 4.3 is essentially the result [12, Proposition 8], which solely relies on the Shapley axioms (Sym)–(Eff). We believe this is evidence that these axioms make more sense for WSMSs. In particular, the axiom (Lin), when instantiated with Ξ^{ms}, implies "if $\mathsf{MS}_{q_1}(\mathcal{D}) \cap \mathsf{MS}_{q_2}(\mathcal{D}) = \varnothing$, then $\phi(\mathcal{D}, q_1, \alpha) + \phi(\mathcal{D}, q_2, \alpha) = \phi(\mathcal{D}, q_1 \wedge q_2, \alpha)$ for all α". This property is key to Proposition 4.3 and further to the efficient computation of the MS Shapley value (see Sections 5 and 6), whereas the instantiation of (Lin) by Ξ^{dr} yields a nonsensical statement.

The first observation that can be made regarding $\phi^{\mathrm{invw}}_{\mathsf{wsms}}$ in contrast with the drastic Shapley value is the fact that only the latter is affected by the way minimal supports may intersect. As such, $\phi^{\mathrm{invw}}_{\mathsf{wsms}}$ will ignore the "green" minimal supports in Example 4.1 and attribute a greater value to e_0. In short the drastic Shapley value considers the interactions between all minimal supports, while WSMSs don't. We find no conceptual reason to argue that one approach gives better explanations than the other. However, we can say that the explanations given by $\phi^{\mathrm{invw}}_{\mathsf{wsms}}$ are *simpler*, both conceptually (it is a straightforward average) and computationally (intuitively, even when the number of minimal supports is manageable, considering all their possible interactions can be costly).

4.3 WSMSs are Shapley-like Scores

In Section 3, we insisted on the term "drastic Shapley value", which is simply called "Shapley value" in the database literature [2, 17]. This is because, as we have seen, the Shapley value proper isn't a quantitative measure, but rather a family of measures parameterized by a wealth function family Ξ. Moreover, while Ξ^{dr} is a very natural choice, it is not the only reasonable wealth function family one can define for non-numeric queries. Indeed, Proposition 4.3 shows that instantiating the Shapley value using Ξ^{ms} yields the intuitive measure $\phi^{\mathrm{invw}}_{\mathsf{wsms}}$. Interestingly, we can generalize the latter result to show that *every* WSMS is actually a Shapley value in disguise:

Proposition 4.4. *For every WSMS* $\phi^{\mathrm{w}}_{\mathsf{wsms}}$ *and every Shapley-like score* $\mathcal{S}_c$ *whose coefficient function* c *is positive on all inputs, there exists a family* $\Xi \doteq \left(\xi_{q,\mathcal{D}_{\mathsf{x}}}\right)$ *such that for every query* q*, database* $\mathcal{D}$ *and* $\alpha \in \mathcal{D}$, $\phi^{\mathrm{w}}_{\mathsf{wsms}}(\mathcal{D}, q, \alpha) = \mathcal{S}_c(\mathcal{D}_{\mathsf{n}}, \xi_{q,\mathcal{D}_{\mathsf{x}}}, \alpha)$.

In light of the previous proposition, given a WSMS ϕ^{w}_{wsms}, we shall denote by Ξ^{w} the wealth function family such that $\phi^{w}_{\mathsf{wsms}}(\alpha) = \mathrm{Sh}(\mathcal{D}_{\mathsf{n}}, \xi^{w}_{q,\mathcal{D}_{\mathsf{x}}}, \alpha)$ and $\mathrm{SVC}^{w}_{\mathcal{C}}$ the associated computational problem. We henceforth abandon the notation ϕ^{w}_{wsms} in favor of the Shapley-based alternative.

4.4 Other Notable WSMSs

By its very nature, the drastic Shapley value gives a lot of importance to small minimal supports, to the point where a fact α such that $\{\alpha\} \models q$ will always receive a maximal drastic Shapley value [2, Lemma 6.3]. This is not the case of the MS Shapley value, which will favor three minimal supports of size 2 against one singleton support. It is however possible to choose a positive and strictly decreasing WSMS such that smaller minimal supports will always dominate.

Proposition 4.5. *Define the weight function $w_s(k,n) \doteq 2^{-kn}$ and denote $\Xi^{s} = \Xi^{w_s}$ for short. The function w_s is tractable, positive and strictly decreasing, and if $\alpha, \beta \in \mathcal{D}_{\mathsf{n}}$ are such that the smallest minimal support containing α is smaller than the smallest minimal support containing β, then $\mathrm{Sh}(\mathcal{D}_{\mathsf{n}}, \xi^{s}_{q,\mathcal{D}_{\mathsf{x}}}, \alpha) > \mathrm{Sh}(\mathcal{D}_{\mathsf{n}}, \xi^{s}_{q,\mathcal{D}_{\mathsf{x}}}, \beta)$, unless $\mathcal{D}_{\mathsf{x}} \models q$.*

Proof sketch. The binary representation of $\mathrm{Sh}(\mathcal{D}_{\mathsf{n}}, \xi^{s}_{q,\mathcal{D}_{\mathsf{x}}}, \alpha)$ will essentially encode the vector of numbers of minimal supports for q that contain α, grouped by increasing size.[9] □

Alternatively, one can also decide that the number of minimal supports should be the main factor for the responsibility measure.

Proposition 4.6. *Define the weight function $w_{\#}(k,n) \doteq 1 + w_s(k,n)$ and denote $\Xi^{\#} = \Xi^{w_{\#}}$ for short. $w_{\#}$ is tractable, positive and strictly decreasing, and if $\alpha, \beta \in \mathcal{D}_{\mathsf{n}}$ are such that α appears in more minimal supports than β, then $\mathrm{Sh}(\mathcal{D}_{\mathsf{n}}, \xi^{\#}_{q,\mathcal{D}_{\mathsf{x}}}, \alpha) > \mathrm{Sh}(\mathcal{D}_{\mathsf{n}}, \xi^{\#}_{q,\mathcal{D}_{\mathsf{x}}}, \beta)$, unless $\mathcal{D}_{\mathsf{x}} \models q$.*

Proof sketch. $\left\lfloor \mathrm{Sh}(\mathcal{D}_{\mathsf{n}}, \xi^{\#}_{q,\mathcal{D}_{\mathsf{x}}}, \alpha) \right\rfloor$ equals the number of minimal supports that contain α. □

5 Data Complexity

We shall now investigate the complexity of computing weighted sums of minimal supports, *i.e.*, computing $\mathrm{SVC}^{w}_{\mathcal{C}}$ for weight functions w. We start with the *data complexity*, *i.e.*, the complexity of SVC^{w}_{q} for individual queries q, independently of their representation.

The simplicity of Equation (3) defining WSMSs not only makes these measures easier to interpret, but also easier to compute in many cases. Essentially, this is because the problem boils down to counting the minimal supports of every size. This straightforward reduction is formalized in the following lemma, which makes use of the aggregate query $\#^{\mathsf{fms}}_{q}$ that, given a number k and database $\mathcal{D}$, computes the number of minimal supports of q in $\mathcal{D}$ having exactly k facts (fms stands for fixed-size ms, in line with 'fixed-size model counting' [13, §3]). For a class $\mathcal{C}$ of Boolean queries, we use $\textsc{eval}\text{-}\#^{\mathsf{fms}}_{\mathcal{C}}$ for the problem of, given a query $q \in \mathcal{C}$, a number k and a database $\mathcal{D}$, evaluating $\#^{\mathsf{fms}}_{q}$ on $k, \mathcal{D}$. If $\mathcal{C}$ contains a single query q, we simply write $\textsc{eval}\text{-}\#^{\mathsf{fms}}_{q}$.

Lemma 5.1. *For a tractable weight function w and Boolean monotone query q, $\mathrm{SVC}^{w}_{q} \leq_{poly} \textsc{eval}\text{-}\#^{fms}_{q}$.*

With Lemma 5.1, we can effectively compute SVC^{w}_{q} for every query whose minimal supports can be enumerated in polynomial time, which includes UCQs or more generally bounded queries.

Theorem 5.2. *$\mathrm{SVC}^{w}_{q} \in \mathsf{FP}$ for every tractable weight function w and every Boolean monotone query q that is bounded and can be evaluated in polynomial time.*

[9] For simplicity this w_s allocates $n = |\mathcal{D}|$ bits to every possible size, but this can obviously be optimized.

Proof. By definition, the minimal supports for q are bounded by some constant M, so one can simply enumerate all subsets of $\mathcal{D}$ of size at most M, and check which ones are minimal supports in order to compute $\#_q^{\mathsf{fms}}(k, \mathcal{D})$ for every k, and finally apply Lemma 5.1. □

Observe that enumerating minimal supports is a somewhat brutal approach to evaluating $\#_q^{\mathsf{fms}}(k, \mathcal{D})$. As exemplified by the following proposition, in some cases counting may be tractable even when the exponential number of minimal supports makes their enumeration intractable.

Proposition 5.3. *Let w be a tractable weight function and q a regular path query. Then, SVC_q^w restricted to acyclic graph databases (i.e., only binary relations and no directed cycles) is in* FP.

Proof sketch. When $\mathcal{D}$ is acyclic, the number of paths of any given size from c to d that are labelled by a word in $\mathcal{L}$ can be computed in polynomial time by dynamic programming. □

On arbitrary graphs containing cycles, deciding whether an edge lies on a simple path is already a known NP-complete problem [7], so no measure that satisfies (Null-db) will be in FP for RPQs, unless P = NP. In fact, we now prove a FP/#P-hard dichotomy (Theorem 5.6) which shows that, for a wide class of WSMSs, their computation is #P-hard for (almost) any unbounded RPQ.

Before we prove this dichotomy, we need to show a lower-bound equivalent of Lemma 5.1. We say a weight function w is *reversible* if the total number of minimal supports for any q in any database $\mathcal{D}$ can be computed in polynomial time given access to an oracle for $\mathsf{Sh}(\mathcal{D}, \xi^w_{q,\emptyset}, \alpha)$, for all $\alpha \in \mathcal{D}$.[10] We shall abuse terminology and say the family Ξ^w itself is reversible when w is.

Remark 5.4. The three families Ξ^{ms}, Ξ^{s} and $\Xi^{\#}$ considered thus far are all reversible.

The next lemma shows how the problem $\textsc{eval}\text{-}\#_q^{\mathsf{ms}}$ of evaluating $\#_q^{\mathsf{ms}}$, i.e., counting, for the given database $\mathcal{D}$, the *total* number of minimal supports of q in $\mathcal{D}$, can yield lower bounds for SVC_q^w.

Lemma 5.5. *For a reversible weight function w and Boolean monotone query q, $\textsc{eval}\text{-}\#_q^{ms} \leq_{poly} \mathsf{SVC}_q^w$.*

We now give the announced dichotomy for RPQs, whose proof relies upon Lemmas 5.1 and 5.5.

Theorem 5.6. *Let w be a reversible tractable weight function, and a regular path query $q \doteq \mathcal{L}(c, d)$. Then $\mathsf{SVC}_q^w \in$ FP if $\mathcal{L}$ is finite or $\varepsilon \in \mathcal{L}$ and $c = d$, and #P-hard otherwise.*

Proof sketch. We reduce from the problem of counting simple paths between two distinguished vertices v_s and v_t in a directed graph, using the pumping lemma for regular languages. □

Note that the proof of Proposition 5.6 fails if we consider existentially quantified RPQs, as discussed in the appendix. However, we have no reason to think that the problem is tractable:

Conjecture 5.7. *Let w be a reversible tractable weight function, and let q be a homomorphism-closed*[11] *Boolean query. Then $\mathsf{SVC}_q^w \in$ FP if q is equivalent to a UCQ and #P-hard otherwise.*

6 Combined Complexity

We will now study the combined complexity of computing weighted sums of minimal supports.

[10]Observe that we do not need to compute the number of minimal supports of each size, but only the total number.
[11]A constant-free query q is *homomorphism-closed* if $\mathcal{D}' \models q$ whenever $\mathcal{D} \xrightarrow{\text{hom}} \mathcal{D}'$ and $\mathcal{D} \models q$, where $\mathcal{D} \xrightarrow{\text{hom}} \mathcal{D}'$ is defined similarly to query homomorphisms but without condition (ii). For queries with constants, we must use the more general notion of C-homomorphism that fixes a set C of constants to ensure that the query constants are preserved [2].

6.1 Arbitrary Query Classes

We first give a general upper bound assuming some minimal conditions, namely that the query evaluation task lies within the polynomial hierarchy and the weight function is tractable.

Proposition 6.1. *Let $\mathcal{C}$ be a class of Boolean monotone queries for which the (combined) query evaluation is in* PH*, and w be a tractable weight function. Then $\mathsf{SVC}^{w}_{\mathcal{C}} \in \mathsf{FP}^{\#\mathsf{P}}$.*

Proof. With Lemma 5.1 in mind, our first step is to show that $\#^{\mathsf{fms}}_{\mathcal{C}} \in \#\mathsf{PH}$, where #PH is the class of all computational problems that can be defined as the number of accepting runs of a polynomial-time nondeterministic Turing machine (PNTM) that has access to a PH oracle [10]. We thus construct a PNTM with PH oracle that, on any input query $q \in \mathcal{C}$ database $\mathcal{D}$, has exactly one accepting run for every minimal support of size k for q in $\mathcal{D}$. For this, we non-deterministically pick a subset S of $\mathcal{D}$ and then use the PH oracle to check if it is a minimal support. Since answering q is in PH, a first call can decide if $S \models q$ and a second one that for all $S' \subseteq S$, $S' \not\models q$.

Now that we have shown $\#^{\mathsf{fms}}_{\mathcal{C}} \in \#\mathsf{PH}$, Lemma 5.5 implies $\mathsf{SVC}^{w}_{\mathcal{C}} \in \mathsf{FP}^{\#\mathsf{PH}}$, and we conclude by the fact that $\mathsf{FP}^{\#\mathsf{PH}} = \mathsf{FP}^{\#\mathsf{P}}$ [27, Theorem 4.1]. □

6.2 (Acyclic) Conjunctive Queries

We will now provide a more detailed analysis for CQs, which will exploit connections to the related problem of counting query homomorphisms. To this end, we define $\#^{\mathsf{hom}}_{q}$ as the query which counts the number of homomorphisms $q \xrightarrow{\text{hom}} \mathcal{D}$ and use $\textsc{eval}\text{-}\#^{\mathsf{hom}}_{\mathcal{C}}$ to refer to the associated evaluation problem for a given class $\mathcal{C}$ of CQs. We observe that for every CQ q and database $\mathcal{D}$, each homomorphism $h : q \xrightarrow{\text{hom}} \mathcal{D}$ *induces* the support $\{h(\alpha) : \alpha \text{ atom of } q\}$, hence $\#^{\mathsf{ms}}_{q}(D) \leqslant \#^{\mathsf{hom}}_{q}(D)$. However, as the next example illustrates, counting homomorphisms is not quite the same as counting minimal supports:

Example 6.2. Consider the CQ $\exists xyzw\ R(x,y) \wedge R(y,z) \wedge R(z,w)$ and the database $\mathcal{D} = \{R(a,b), R(b,c), R(c,a)\}$. While the query has only one (minimal) support (*i.e.*, $\mathcal{D}$ itself), it has three different homomorphisms characterized by their values on x, y, z, w: $\{(a,b,c,a), (b,c,a,b), (c,a,b,c)\}$. △

It is known that counting homomorphisms from CQs (or, equivalently, counting the answers to non-Boolean CQs without any existential variable) is #P-complete in general [6, Theorem 1.1], but it becomes tractable for ACQs [22, Theorem 1] and more generally for classes of CQs having bounded 'generalized hypertree width' [22, Theorem 7]. Interestingly, we show here that these tractability results do not generalize to counting minimal supports. In fact, we prove that counting minimal supports for ACQs is #P-hard even in the absence of constants and over a fixed schema.

Theorem 6.3. *$\textsc{eval}\text{-}\#^{ms}_{\mathsf{ACQ}}$ is #P-hard under polynomial-time 1-Turing reductions[12], even in the absence of constants. Hence, $\textsc{eval}\text{-}\#^{fms}_{\mathsf{ACQ}}$ is #P-hard under polynomial-time Turing reductions.*

Proof. We reduce from the *perfect matching counting* problem, which is the problem of, given a bipartite graph $G = (V, E)$, counting the number of subsets $M \subseteq E$ (called *perfect matchings*) such that each vertex of V is contained in exactly one edge of M. This problem is known to be #P-complete [28, Problem 2] under polynomial-time Turing reductions, and was later shown to be hard even under polynomial-time 1-Turing reductions [29].

Suppose V is partitioned into $V_1 \cup V_2$ and assume $V_1 = \{a_1, \ldots, a_n\}$, $V_2 = \{b_1, \ldots, b_n\}$ and denote $I = \{1, \ldots, n\}$ (observe that if $|V_1| \neq |V_2|$ there are no perfect matchings). We will build a database D_G and a constant-free ACQ q_G such that the number of perfect matchings on G is equal to

[12]That is, Turing reductions that only allow a single call to an oracle.

$\#^{\mathsf{ms}}_{q_G}(D_G) - n$. For clarity, we first show a reduction which uses an unbounded relational schema $\sigma_G = \{R_1, \dots, R_n\}$ where each R_i is binary.

D_G We shall abuse notation and use the elements of $V_1 \cup V_2$ as constant names. The database D_G is built over the domain $V_1 \cup V_2 \cup \{c_{i,v,v'} \mid i \in [n-1], v, v' \in V_1 \cup V_2\}$. The c-constants are there only to build some paths between the elements of $V_1 \cup V_2$ reading '$R_1 \cdots R_n$' or '$R_n \cdots R_1$'. Figure 4 contains an example from which the reader can deduce how D_G is defined (the c-constants are implicit in the figure). More formally, we define D_G as the smallest set containing:

- All facts $R_1(v,v), \dots, R_n(v,v)$ for every $v \in V_2$.
- A path reading $R_n, R_{n-1}, \dots, R_1$ from a_i to a_{i+1} for every i —formally, all facts $R_n(a_i, c_{1,a_i,a_{i+1}})$, $R_{n-1}(c_{1,a_i,a_{i+1}}, c_{2,a_i,a_{i+1}}), \dots, R_1(c_{n-1,a_i,a_{i+1}}, a_{i+1})$ for every $i \in [n-1]$.
- A path reading $R_1, R_2, \dots, R_n$ between any a_i and b_j connected by an edge in G —formally, all facts $R_1(a_i, c_{1,a_i,b_j}), R_2(c_{1,a_i,b_j}, c_{2,a_i,b_j}), \dots, R_n(c_{n-1,a_i,b_j}, b_j)$ for every $\{a_i, b_j\} \in E$.

q_G The ACQ q_G is defined as in Figure 4. We give names only to some of the variables of q_G: $X = \{x_1, \dots, x_n\}$ which will be mapped to the a_i's, and $Y = \{y_1, \dots, y_n\}$ which will be mapped to the b_i's. All others will be mapped to intermediate constants along paths in $\mathcal{D}_G$.

It is easy to see that the $\mathcal{S}_j \mathrel{\hat{=}} R_I(b_j, b_j)$ (*i.e.* $R_1(b_j,b_j), \dots, R_n(b_j,b_j)$) are all minimal supports for q_G in $\mathcal{D}_G$. Similarly, if we take a perfect matching in G under the form of a bijective $\mu : V_1 \to V_2$ *s.t.* $\forall j \in I.(a_j, \mu(a_j)) \in E$, then we can easily build a minimal support for q_G by taking an isomorphic image of q_G with $(x_j, y_j) \mapsto (a_j, \mu(a_j))$. We now show that there is no minimal support for q_G in $\mathcal{D}_G$ other than these.

Assume that $\mathcal{S}$ is a minimal support for q_G in $\mathcal{D}_G$ that doesn't contain any $\mathcal{S}_j$. Since $\mathcal{S} \models q$, then $\mathcal{S}$ must be the homomorphic image of q_G by some homomorphism h. First consider the spine $x_1 \xrightarrow{R_{n\dots1}} \dots \xrightarrow{R_{n\dots1}} x_n$ of q_G. It must necessarily map to $a_1 \xrightarrow{R_{n\dots1}} a_n$ because that is the only path in $\mathcal{D}_G$ that doesn't contain any $\mathcal{S}_j$ and is homomorphic to $x_1 \xrightarrow{R_{n\dots1}} \dots \xrightarrow{R_{n\dots1}} x_n$. In other words $\forall j \in I.h(x_j) = a_j$.

Now consider a branch $x_i \xrightarrow{R_{1\dots n}} y_i$ of q_G. Since $h(x_i) = a_i$ and x_i has at most two outgoing edge, only one of which is a R_1, the branch must necessarily map to some $a_i \xrightarrow{R_{1\dots n}} b_j$, with $(a_i, b_j) \in E$. In other words, $h(y_j)$ is always a neighbour of a_j in G. Moreover, if we were to have $h(y_i) = h(y_{i'}) = b_j$ with $i \neq i'$, then $\mathcal{S}$ would necessarily contain $R_{I\setminus\{i\}}(b_j,b_j) \cup R_{I\setminus\{i'\}}(b_j,b_j) = \mathcal{S}_j$, which is impossible. Overall, $\mathcal{S}$ must therefore induce a perfect matching as described above, with $\mu(a_j) \mathrel{\hat{=}} h(y_j)$.

One shortcoming of the reduction above is that we have used a schema which depends on the problem instance. However, this was done for the sake of clarity of exposition, and we show how to modify the reduction above to use a *bounded* schema $\sigma = \{S, T, R\}$, where S is unary and T, R are binary. For this, we replace everywhere (in the query and in the database) each appearance of $R_i(t, t')$ with a path from t to t' reading TR^i (by using fresh i new constants or variables). The proof follows the same lines. □

Corollary 6.4. $\mathrm{SVC}^{w}_{\mathrm{ACQ}}$ *is #P-hard for any reversible weight function* w. *In particular for* $\mathrm{SVC}^{ms}_{\mathrm{ACQ}}$.

Proof. This is a direct consequence of Lemma 5.5, Theorem 6.3, and Remark 5.4. □

6.3 Conjunctive Queries with Bounded Self-joins

In order to identify some tractable cases, we will turn our attention to the class sjf-CQ of *self-join free* CQs, which contains CQs which do not contain two distinct atoms with the same relation name. The interest of this class lies in the fact that for self-join free CQs, the number of minimal supports always coincides with the number of homomorphisms:

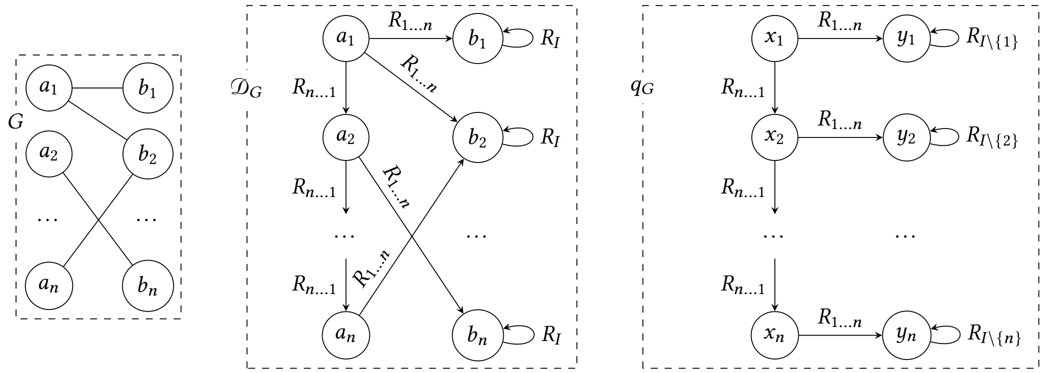

Fig. 4. Construction for the reduction in the proof of Theorem 6.3. $\xrightarrow{R_{1\ldots n}}$ (resp. $\xrightarrow{R_{n\ldots 1}}$) is an abbreviation for the path $\xrightarrow{R_1} \ldots \xrightarrow{R_n}$ (resp. $\xrightarrow{R_n} \ldots \xrightarrow{R_1}$), and $\xrightarrow{R_J}$ means a parallel edge for every $j \in J$.

Remark 6.5. For every self-join free CQ q and database $\mathcal{D}$, $\#^{\mathsf{hom}}_q(\mathcal{D}) = \#^{\mathsf{ms}}_q(\mathcal{D}) = \#^{\mathsf{fms}}_q(k, \mathcal{D})$, where $k = |atoms(q)|$ (in particular, $\#^{\mathsf{fms}}_q(k', \mathcal{D}) = 0$ for any $k' \neq k$).

As a corollary, we have that $\textsc{eval}\text{-}\#^{\mathsf{fms}}_{\mathsf{sjf\text{-}ACQ}}$ is tractable for the class sjf-ACQ of acyclic self-join free CQs since it is tractable for $\textsc{eval}\text{-}\#^{\mathsf{hom}}_{\mathsf{sjf\text{-}ACQ}}$ [22, Theorem 1]. This means that we can reuse the efficient divide-and-conquer algorithms exploiting the tree-like structure of queries that have been previously devised for counting homomorphisms.

In the presence of self-joins, we know from Theorem 6.3 that the problem remains #P-hard, but the proof uses an *unbounded* number of self-joins, raising the obvious question of what happens if the number of self-joins is bounded. We shall now show a positive result, namely, that counting minimal supports is in polynomial time if we simultaneously bound the number of self-joins and restrict to acyclic or more generally tree-like conjunctive queries. The measure of "tree-like" which we will use here is that of *generalized hypertree width*, which we will not define here but can be found, *e.g.*, in [8, Definition 3.1]. ACQs correspond to CQs of generalized hypertree width 1.

Two distinct atoms α, β of a CQ q are *mergeable* if there are two homomorphisms $h_\alpha : \alpha \xrightarrow{\text{hom}} \mathcal{D}$, $h_\beta : \beta \xrightarrow{\text{hom}} \mathcal{D}$ to an arbitrary database $\mathcal{D}$ such that $h_\alpha(\alpha) = h_\beta(\beta)$ (in particular they must have the same relation name). An atom α is (individually) mergeable if it is mergeable with some other atom in q. The *self-join width* of a CQ q is the size of

$$\mathbf{M}_q \doteq \{t \in \text{terms}(\alpha) : \alpha \text{ is a mergeable atom of } q\}.$$

In particular, self-join free queries have self-join width 0. Note that a CQ like $q \doteq \exists xy R(c, x) \wedge R(c', y)$, where c, c' are distinct constants, also has self-join width 0 since it has no two mergeable atoms. Also observe that in the absence of mergeable atoms, each minimal support corresponds to a distinct homomorphism, and each homomorphism gives rise to a distinct minimal support.

Theorem 6.6. *For any class $\mathcal{C}$ of CQs having bounded generalized hypertree width and bounded self-join width, $\textsc{eval}\text{-}\#^{fms}_{\mathcal{C}}$ and $\textsc{eval}\text{-}\#^{ms}_{\mathcal{C}}$ are in polynomial time.*[13]

Proof sketch. We will use CQs with equality and/or disequality atoms denoted by $CQ^{\neq}$, $CQ^{\neq,=}$, $CQ^{=}$, where the notions of homomorphisms from such queries to databases is restricted accordingly.

[13] We understand this statement as a "promise problem", that is, the input to the evaluation problem is a query from $\mathcal{C}$, as opposed to some arbitrary string.

Observe that any equality atom $x = t$ can always be removed by replacing every occurrence of x with t (or removing the equality if $t = x$). We write $\widehat{q}$ to denote the canonical[14] equivalent $\mathsf{CQ}^{\neq}$ (resp. CQ) query resulting from applying the preceding simplifications, for any $\mathsf{CQ}^{\neq,=}$ (resp. $\mathsf{CQ}^{=}$) q.

Let q be a CQ from $\mathcal{C}$. Consider $\mathcal{E}$ to be the set of all *equivalence relations* over $\mathbf{M}_q$ (*i.e.*, relations $E \subseteq \mathbf{M}_q \times \mathbf{M}_q$ being reflexive, symmetric and transitive) not containing any pair (c, c') of distinct constants. For any $E \in \mathcal{E}$, define $q_E \in \mathsf{CQ}$ as $core(\widehat{p})$, where $p \in \mathsf{CQ}^{=}$ is defined as the result of adding to the query q the equality atoms $t = t'$ for every $(t, t') \in E$. Let $q_E^{\neq} \in \mathsf{CQ}^{\neq}$ be q_E with some added disequality atoms $t \neq t'$ for $t, t' \in \mathrm{terms}(q_E)$ such that $(t, t') \in (\mathbf{M}_q \times \mathbf{M}_q) \setminus E$. Let $\mathbf{P}_q \mathrel{\hat{=}} \{q_E : E \in \mathcal{E} \text{ s.t. } q_{E'}^{\neq} \xrightarrow{\text{hom}} q_E^{\neq} \text{ implies } q_E^{\neq} \xrightarrow{\text{hom}} q_{E'}^{\neq} \text{ for all } E' \in \mathcal{E}\}$. The proof of the statement can be shown by establishing the following facts:

(1) If a homomorphism $h : q_E^{\neq} \xrightarrow{\text{hom}} \mathcal{D}$ induces S, then $|S|$ is the number of atoms of q_E.
(2) $\mathbf{P}_q$ can be computed in polynomial time.
(3) The set of minimal supports of q on $\mathcal{D}$ is partitioned into

$$\{S \subseteq \mathcal{D} : S \text{ is induced by a homomorphism } q_E^{\neq} \xrightarrow{\text{hom}} \mathcal{D}\}_{q_E \in \mathbf{P}_q}.$$

(4) For every $q_E \in \mathbf{P}_q$ and database $\mathcal{D}$, we have

$$\#^{\mathsf{hom}}_{q_E^{\neq}}(\mathcal{D}) = |\mathrm{Aut}(q_E^{\neq})| \cdot |\{S : S \text{ induced by homomorphism } q_E^{\neq} \xrightarrow{\text{hom}} \mathcal{D}\}|$$

where $\mathrm{Aut}(q_E^{\neq})$ is the set of automorphisms $q_E^{\neq} \xrightarrow{\text{hom}} q_E^{\neq}$.
(5) $|\mathrm{Aut}(q_E^{\neq})|$ and $\#^{\mathsf{hom}}_{q_E^{\neq}}(\mathcal{D})$ can be computed in polynomial time.

Hence, for every fixed k, $\#^{\mathsf{fms}}_q(k, \mathcal{D})$ can be computed in polynomial time: it suffices to compute, in polynomial time, $\sum_{q_E \in P_k} (\#^{\mathsf{hom}}_{q_E^{\neq}}(\mathcal{D}) / |\mathrm{Aut}(q_E^{\neq})|)$ for the set $P_k \subseteq \mathbf{P}_q$ of queries having k atoms. □

7 Related Measures

We review some measures known in the literature and compare them with the WSMS family.

7.1 Shapley Value of Inconsistency Measures

As discussed earlier, the drastic and the MS Shapley values are analogues to previously studied inconsistency measures [12, 18]. Livshits and Kimelfeld [18] have considered three further inconsistency measures, denoted by I_{P}, I_{MC} and I_{R} respectively, defined in the absence of exogenous atoms. The first two translate into the following families of wealth functions:

- Ξ^{P} is the analogue of I_{P}: $\xi^{\mathsf{P}}_{q,\varnothing}(S)$ is the number of facts that are relevant to q in S.
- Ξ^{MC} is the analogue of I_{MC}: $\xi^{\mathsf{MC}}_{q,\varnothing}(S) = mc(S) + ss(S) - 1$ where $mc(S)$ is the number of maximal subsets $\Delta \subseteq S$ s.t. $\Delta \not\models q$, and $ss(S) \mathrel{\hat{=}} |\{s \in S : \{s\} \models q\}|$.[15]

We did not extend these definitions to add exogenous facts, because the following result proves that they are not suitable responsibility measures using our criteria, even without them.

Proposition 7.1. *The Shapley values of Ξ^{P} and Ξ^{MC} don't satisfy (MStest), even assuming $\mathcal{D}_{\mathsf{x}} = \varnothing$.*

Proof. The Shapley value of Ξ^{P} attributes the score 1 to all facts in the example depicted in Figure 1-(a) ($k = l = 1, |A_1| = 1, |B_1| = 2$ with the notations of (MStest)), and the one of Ξ^{MC} attributes the score $144/120$ to all facts in the variant defined by $k = l = 1, |A_1| = 2, |B_1| = 3$. □

We do extend the definition of the third to account for $\mathcal{D}_{\mathsf{x}}$, in order to prove its suitability.

[14] Canonical in the sense of unique up to variable renaming.
[15] The definition in [18] is simpler than ours because in their context (functional dependencies), the analogue of such a "singleton support" cannot appear, hence $ss(S) = 0$; the full definition can be found in [9, Definition 2] (I_M measure).

- Ξ^{R} is the analogue of I_{R}: $\xi^{\mathsf{R}}_{q,\mathcal{D}_{\mathsf{x}}}(S)$ is the size of the smallest $\Gamma \subseteq S$ such that $(S \setminus \Gamma) \cup \mathcal{D}_{\mathsf{x}} \not\models q$.

Proposition 7.2. *The Shapley value of Ξ^R satisfies (Sym-db), (Null-db) and (MStest).*

These results show that Ξ^{P} and Ξ^{MC} do not yield suitable responsibility measures, but Ξ^{R} remains a reasonable option. While we shall leave a more thorough investigation of Ξ^{R} to future work (in particular, it remains to determine if the Shapley value of Ξ^{R} satisfies (MS1) and (MS2) as formally defined in Appendix B.1), we can already show that it is difficult to compute. Indeed, not only is Ξ^{R} intractable in data complexity (due to the similarity with Ξ^{dr}) but it is not even approximable via a fully polynomial randomized approximation scheme (FPRAS).

Proposition 7.3. *There exist CQs q, q' such that SVC^R_q is #P-hard and such that $\mathsf{SVC}^R_{q'}$ possesses no FPRAS unless $\mathsf{NP} \subseteq \mathsf{BPP}$ nor any multiplicative approximation of factor 1.3606 unless $\mathsf{P} = \mathsf{NP}$.*

7.2 Counting Homomorphisms of Minimal Supports

A natural idea for an alternative measure for (Boolean) UCQs would be to replace "counting minimal supports" with "counting homomorphisms" in the definition of ξ^{ms} to obtain a wealth function family Ξ^{hom} based on counting homomorphisms (*i.e.*, $\xi^{\mathsf{hom}}_{q,\mathcal{D}_{\mathsf{x}}}(\mathcal{D}_{\mathsf{n}}) \mathrel{\hat{=}} |\{h : q \xrightarrow{\mathsf{hom}} \mathcal{D}_{\mathsf{n}}\}|$). Adopting Ξ^{hom} in place of Ξ^{ms} would allow us to obtain further tractability results, as we have seen that counting homomorphisms can be computationally simpler than counting minimal supports (Theorem 6.3). Unfortunately, although Ξ^{hom} seems reasonable at first sight, we show that it cannot decide relevance, which may be considered a minimum requirement for any responsibility measure:

Proposition 7.4. *The Shapley value of Ξ^{hom} does not satisfy (Null-db). Moreover, an irrelevant fact may have a larger value than relevant ones by an arbitrarily large (multiplicative or additive) margin.*

Proof sketch. Consider the CQ $q \mathrel{\hat{=}} \exists xyz\ R(x,y) \wedge R(y,z)$ and the purely endogenous database $\mathcal{D} \mathrel{\hat{=}} \{R(c,c), R(c,d)\}$. The fact $R(c,d)$ appears in the assignment $(c,c,d) \mapsto (x,y,z)$ hence $\mathsf{Sh}(\mathcal{D}, \xi^{\mathsf{hom}}_{q,\varnothing}, R(c,d)) > 0$, but it is irrelevant because $R(c,c)$ satisfies q on its own. A slight variation of the above instance can get an irrelevant fact to have the highest value by an arbitrary margin. □

7.3 Explainable AI and SHAP Score

The Shapley value has also been applied to define a popular quantitative measure of responsibility for machine learning, namely the SHAP score [19], which we shall briefly define later. Despite its popularity in practice, recent work [20] has cast doubts on the theoretical guarantees this measure provides. Although the context is substantially different from query answering, we discuss it here because we believe the present work can provide valuable insights on the issue.

Consider a *classifier*, which we shall simply model as a function $\mathcal{M} : \mathbb{B}^{\mathcal{F}} \to \mathbb{B}$, where $\mathcal{F}$ is a finite set of *features* and $\mathbb{B}$ is the Boolean set $\{0,1\}$.[16] Given an input instance $\alpha \in \mathbb{B}^{\mathcal{F}}$, we wish to explain the classification $\mathcal{M}(\alpha)$ by identifying the most relevant features.

In order to obtain a quantitative score from the Shapley value, one simply needs to define a relevant game whose players are the features in $\mathcal{F}$. In the context of query answering, we have seen that one straightforward approach (giving rise to the drastic Shapley value) consists in using the query itself, seen as a 0/1 function over the subsets of the input. But this is not possible for classifiers because $\mathcal{M}$ needs all the features of its input to have some value. The solution chosen to define the SHAP score consists in taking the average value of $\mathcal{M}$ across all completions of $\alpha|_S$ (that is, α restricted to the features in S): $\xi^{\mathsf{SHAP}}(S) \mathrel{\hat{=}} \frac{1}{2^{|\mathcal{F} \setminus S|}} \sum_{x \in \mathbb{B}^{\mathcal{F}},\, x|_S = \alpha|_S} \mathcal{M}(x)$. The *SHAP score* is then defined as the Shapley value applied to the game $(\mathcal{F}, \xi^{\mathsf{SHAP}})$.

[16]For the sake of simplicity, we only consider Boolean classifiers here.

In order to analyse the SHAP score, [20] defines the notion of an *abductive explanation* (AXp for short) as an inclusion-minimal subset $S \subseteq \mathcal{F}$ such that for all $x \in \mathbb{B}^{\mathcal{F}}$, $x|_S = \alpha|_S \Rightarrow \mathcal{M}(x) = \mathcal{M}(\alpha)$. AXps are precisely analogous to the minimal supports of the database setting, and [20] essentially shows that the SHAP score does not satisfy the analogue of the (Null-db) axiom, which we argued was a very natural condition in our setting.

However, due to the closeness of AXps and minimal supports, one can easily adapt Section 4 of the present paper to define 'Weighted Sums of AXps' (*WSAXp*s). These quantitative measures, which also derive from the Shapley value, will naturally display all of the good properties of WSMSs and therefore make viable candidates for addressing the theoretical drawbacks of the SHAP score. Unfortunately, however, whereas minimal supports can be efficiently counted for relevant classes of database queries, the problem of counting the number of the AXps of each size will be intractable for pretty much any interesting class of classifiers. Moreover, this high complexity is inescapable: as deciding feature relevance is NP-hard for very simple classes of classifiers [11, Proposition 11], no measure satisfying the ML analogue of (Null-db) can be tractable.

8 Conclusions

We have revisited the question of how the Shapley value can be used to define responsibility measures for non-numeric queries. Our proposed family of responsibility measures – weighted sums of minimal supports (WSMSs) – can be viewed as the Shapley values of suitably defined games and have been shown to satisfy desirable semantic properties. Moreover, due to their simple definition, WSMSs are amenable to interpretation and can be efficiently computed for a large class of queries (including all UCQs), making them an appealing alternative to the drastic Shapley value. We should point out that the two main measures that have been studied in the literature, namely the drastic Shapley value and the *drastic Banzhaf value* (*i.e.*, the Banzhaf power index applied to the drastic wealth function family) both satisfy our new desiderata for responsibility measures. More precisely, they satisfy the axioms (Sym-db), (Null-db) and (MStest) by Proposition 4.2, as well as (the formal versions of) the stronger (MS1) and (MS2) by Propositions B.1 and B.2 in the appendix. Therefore, both remain perfectly adequate and sensible measures despite the shortcomings discussed.

Our work has focused on monotone queries, but the drastic Shapley value has also been applied to CQs with negative atoms [23]. A natural next step would thus be to explore alternative responsibility measures for non-monotone queries, in particular, a suitable adaptation of WSMSs. In the presence of exogenous facts, the notion of *generalized supports*, which is very close to that of supports, has proven quite relevant [2]. Replacing minimal supports by minimal generalized supports in the definition of WSMSs could therefore provide an interesting alternative framework in this context. We have also suggested WSAXps as alternatives to the SHAP score in the context of XAI. Although WSAXps are computationally expensive, they provide strong theoretical guarantees, so it would be interesting to study approximations and to compare them experimentally with the SHAP score.

Acknowledgments

All authors have been partially supported by ANR AI Chair INTENDED, grant ANR-19-CHIA-0014.

References

[1] Omer Abramovich, Daniel Deutch, Nave Frost, Ahmet Kara, and Dan Olteanu. 2024. Banzhaf Values for Facts in Query Answering. *Proceedings of the ACM on Management of Data (PACMMOD)* 2, 3 (SIGMOD) (2024), 123:1–123:26. doi:10.1145/3654926

[2] Meghyn Bienvenu, Diego Figueira, and Pierre Lafourcade. 2024. When is Shapley Value Computation a Matter of Counting? *Proceedings of the ACM on Management of Data (PACMMOD)* 2, 2 (PODS) (2024), 105. doi:10.1145/3651606

[3] Ashok K. Chandra and Philip M. Merlin. 1977. Optimal Implementation of Conjunctive Queries in Relational Data Bases. In *Proceedings of the 9th Annual ACM Symposium on Theory of Computing, May 4-6, 1977, Boulder, Colorado,*

USA, John E. Hopcroft, Emily P. Friedman, and Michael A. Harrison (Eds.). ACM, 77–90. doi:10.1145/800105.803397
[4] Daniel Deutch, Nave Frost, Benny Kimelfeld, and Mikaël Monet. 2022. Computing the Shapley Value of Facts in Query Answering. In *ACM SIGMOD International Conference on Management of Data (SIGMOD)*. ACM, 1570–1583. doi:10.1145/3514221.3517912
[5] P. Dubey. 1975. On the Uniqueness of the Shapley Value. *International Journal of Game Theory* 4, 3 (1975), 131–139. doi:10.1007/BF01780630
[6] Martin E. Dyer and Catherine S. Greenhill. 2000. The complexity of counting graph homomorphisms. *Random Struct. Algorithms* 17, 3-4 (2000), 260–289. doi:10.1002/1098-2418(200010/12)17:3/4<260::AID-RSA5>3.0.CO;2-W
[7] Steven Fortune, John Hopcroft, and James Wyllie. 1980. The Directed Subgraph Homeomorphism Problem. *Theoretical Computer Science* 10, 2 (1980), 111–121. doi:10.1016/0304-3975(80)90009-2
[8] Georg Gottlob, Gianluigi Greco, Nicola Leone, and Francesco Scarcello. 2016. Hypertree Decompositions: Questions and Answers. In *ACM Symposium on Principles of Database Systems (PODS)*. ACM Press, 57–74. https://doi.org/10.1145/2902251.2902309
[9] John Grant and Anthony Hunter. 2011. Measuring Consistency Gain and Information Loss in Stepwise Inconsistency Resolution. In *Proceedings of the 11th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'11)*. Springer-Verlag, Berlin, Heidelberg, 362–373.
[10] Lane A. Hemaspaandra and Heribert Vollmer. 1995. The Satanic Notations: Counting Classes beyond #P and Other Definitional Adventures. *ACM SIGACT News* 26, 1 (1995), 2–13. doi:10.1145/203610.203611
[11] Xuanxiang Huang, Martin C. Cooper, Antonio Morgado, Jordi Planes, and Joao Marques-Silva. 2023. Feature Necessity & Relevancy in ML Classifier Explanations. In *Tools and Algorithms for the Construction and Analysis of Systems*, Sriram Sankaranarayanan and Natasha Sharygina (Eds.). Springer Nature Switzerland, Cham, 167–186. doi:10.1007/978-3-031-30823-9_9
[12] Anthony Hunter and Sébastien Konieczny. 2010. On the Measure of Conflicts: Shapley Inconsistency Values. *Artificial Intelligence* 174, 14 (2010), 1007–1026. doi:10.1016/j.artint.2010.06.001
[13] Ahmet Kara, Dan Olteanu, and Dan Suciu. 2024. From Shapley Value to Model Counting and Back. *Proceedings of the ACM on Management of Data (PACMMOD)* 2, 2 (PODS) (2024), 79. doi:10.1145/3651142
[14] Pratik Karmakar, Mikaël Monet, Pierre Senellart, and Stephane Bressan. 2024. Expected Shapley-Like Scores of Boolean Functions: Complexity and Applications to Probabilistic Databases. *Proc. ACM Manag. Data* 2, 2 (2024), 92:1–92:26. doi:10.1145/3651593
[15] Majd Khalil and Benny Kimelfeld. 2023. The Complexity of the Shapley Value for Regular Path Queries. In *International Conference on Database Theory (ICDT) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 255)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 11:1–11:19. doi:10.4230/LIPICS.ICDT.2023.11
[16] Annick Laruelle and Federico Valenciano. 2001. Shapley-Shubik and Banzhaf Indices Revisited. *Mathematics of Operations Research* (2001). doi:10.1287/moor.26.1.89.10589
[17] Ester Livshits, Leopoldo E. Bertossi, Benny Kimelfeld, and Moshe Sebag. 2021. The Shapley Value of Tuples in Query Answering. *Logical Methods in Computer Science (LMCS)* 17, 3 (2021). doi:10.46298/LMCS-17(3:22)2021
[18] Ester Livshits and Benny Kimelfeld. 2022. The Shapley Value of Inconsistency Measures for Functional Dependencies. *Logical Methods in Computer Science (LMCS)* 18, 2 (2022). doi:10.46298/LMCS-18(2:20)2022
[19] Scott M Lundberg and Su-In Lee. 2017. A Unified Approach to Interpreting Model Predictions. In *Advances in Neural Information Processing Systems*, Vol. 30. Curran Associates, Inc.
[20] Joao Marques-Silva and Xuanxiang Huang. 2024. Explainability Is Not a Game. *Commun. ACM* 67, 7 (2024), 66–75. doi:10.1145/3635301
[21] Alexandra Meliou, Wolfgang Gatterbauer, Katherine F. Moore, and Dan Suciu. 2010. The Complexity of Causality and Responsibility for Query Answers and non-Answers. *Proc. VLDB Endow.* 4, 1 (2010), 34–45.
[22] Reinhard Pichler and Sebastian Skritek. 2013. Tractable counting of the answers to conjunctive queries. *Journal of Computer and System Sciences (JCSS)* 79, 6 (2013), 984–1001. doi:10.1016/J.JCSS.2013.01.012
[23] Alon Reshef, Benny Kimelfeld, and Ester Livshits. 2020. The Impact of Negation on the Complexity of the Shapley Value in Conjunctive Queries. In *ACM Symposium on Principles of Database Systems (PODS)*. ACM, 285–297. doi:10.1145/3375395.3387664
[24] Babak Salimi, Leopoldo Bertossi, Dan Suciu, and Guy Van den Broeck. 2016. Quantifying Causal Effects on Query Answering in Databases. In *8th USENIX Workshop on the Theory and Practice of Provenance (TaPP)*.
[25] Lloyd S Shapley. 1953. A value for n-person games. *Contribution to the Theory of Games* 2 (1953).
[26] Matthias Thimm. 2017. On the Compliance of Rationality Postulates for Inconsistency Measures: A More or Less Complete Picture. *Künstliche Intell.* 31, 1 (2017), 31–39. doi:10.1007/S13218-016-0451-Y
[27] Seinosuke Toda and Osamu Watanabe. 1992. Polynomial-Time 1-Turing Reductions from #PH to #P. *Theoretical Computer Science* 100, 1 (1992), 205–221. doi:10.1016/0304-3975(92)90369-Q

[28] Leslie G. Valiant. 1979. The Complexity of Enumeration and Reliability Problems. *SIAM J. Comput.* 8, 3 (1979), 410–421. doi:10.1137/0208032

[29] Viktória Zankó. 1991. #P-Completeness via Many-One Reductions. *Int. J. Found. Comput. Sci.* 2, 1 (1991), 77–82. doi:10.1142/S0129054191000066

Received December 2024; revised February 2025; accepted March 2025

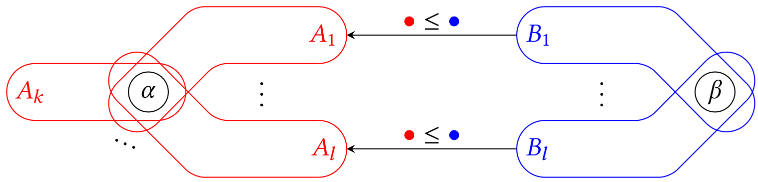

Fig. 3. A visual representation of the setting considered by (MStest). The A_i, B_j are the minimal supports of q.

A Appendix to Section 3 "Responsibility Measures and the Drastic Shapley Value"

Remark A.1. The following alternative to the linearity axiom that still makes sense for Boolean games has been introduced in [5, Axiom S3'], and later used in more modern axiomatizations of the Shapley value restricted to monotone Boolean games, also called Shapley-Shubik value [16, §2]:

(Tr) *Transfer:* Given two wealth functions ξ_1 and ξ_2 with values in $\mathbb{B}$,

$$\psi(P, \xi_1 \vee \xi_2, \alpha) + \psi(P, \xi_1 \wedge \xi_2, \alpha) = \psi(P, \xi_1, \alpha) + \psi(P, \xi_2, \alpha)$$

where $\vee$ and $\wedge$ are the Boolean disjunction and conjunction respectively. Since this axiom is by design restricted to Boolean wealth functions, it can be instantiated with the Ξ^{dr} family:

(Tr-dr) *Transfer:* Given two Boolean queries q_1 and q_2 with values in $\mathbb{B}$,

$$\phi(\mathcal{D}, q_1 \vee q_2, \alpha) + \phi(\mathcal{D}, q_1 \wedge q_2, \alpha) = \phi(\mathcal{D}, q_1, \alpha) + \phi(\mathcal{D}, q_2, \alpha)$$

Although (Tr) is a nice axiom in that it allows to show the unicity of the Shapley value restricted to Boolean games, we don't believe (Tr-dr) to be an obviously desirable feature for a responsibility measure. Even if one were to believe so, it should be noted that all WSMSs satisfy it, by virtue of the fact that they are Shapley values in disguise (see Proposition 4.4).

B Appendix to Section 4 "Alternative responsibility measures"

B.1 Formal statements of (MS1) and (MS2)

In this section we present and explain formal definitions of the axioms (MS1) and (MS2). Recall that they were informally written as follows.

(MS1) All other things being equal, a fact that appears in larger minimal supports should have lower responsibility.

(MS2) All other things being equal, a fact that appears in fewer minimal supports should have lower responsibility.

The key to formalizing these axioms is to define a good notion of *domination witness* $\Delta(\alpha, \beta)$ that should ensure that α must dominate β. A first attempt simply relies on the canonical way of comparing the cardinals of two sets, which consists in finding some injective mapping f between them. Given a fact $\alpha \in \mathcal{D}$, we say that $S \subseteq \mathcal{D} \setminus \{\alpha\}$ is a *completing set* for α if $S \cup \{\alpha\} \in \mathrm{MS}_q(\mathcal{D})$, and denote the set of all completing sets for α by $\mathrm{CS}(\mathcal{D}, q, \alpha)$.

An *a-domination witness* $\Delta(\alpha, \beta)$ of α over β is given by:

- an injective mapping $f : \mathrm{CS}(\mathcal{D}, q, \beta) \hookrightarrow \mathrm{CS}(\mathcal{D}, q, \alpha)$;
- for every $S \in \mathrm{CS}(\mathcal{D}, q, \beta)$, an injective mapping $\eta_S : f(S) \hookrightarrow S$.

(MS1a) If there exists an a-domination witness $\Delta(\alpha, \beta)$ such that one of the $\eta_S : f(S) \hookrightarrow S$ isn't surjective, then $\phi(\mathcal{D}, q, \alpha) > \phi(\mathcal{D}, q, \beta)$.

(MS2a) If there exists an a-domination witness $\Delta(\alpha, \beta)$ such that $f : \mathrm{CS}(\mathcal{D}, q, \beta) \hookrightarrow \mathrm{CS}(\mathcal{D}, q, \alpha)$ isn't surjective, then $\phi(\mathcal{D}, q, \alpha) > \phi(\mathcal{D}, q, \beta)$.

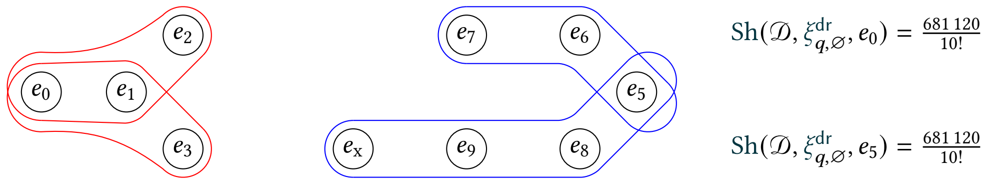

Fig. 6. Example of intersecting minimal supports. The database $\mathcal{D}$ is the set $\{e_0 \dots e_9, e_x\}$, and the query q is such that the minimal supports are the colored regions (the red ones contain e_0 and the blue ones contain e_5).

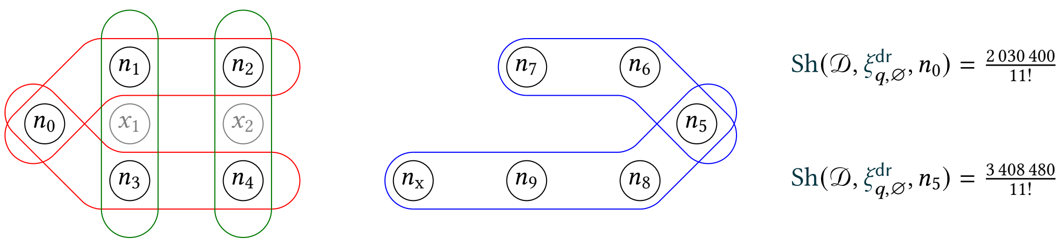

Fig. 2. Instance for Example 4.1. $\mathcal{D}_n \mathrel{\hat{=}} \{n_1, \dots, n_x\}$, $\mathcal{D}_x \mathrel{\hat{=}} \{x_1, x_2\}$, and the query q is such that the minimal supports are the colored regions (the red ones contain e_0, the blue ones e_5 and the green ones neither).

Observe that in the simplified setting used to define (MStest) and depicted in Figure 3, f corresponds to the black arrows and each η_S certifies that $|f(S)| \leq |S|$.

However, the notion of a-domination witness is too simplistic because it considers all minimal supports as independent when in reality they can share many elements. To illustrate this consider the example represented in Figure 6. e_0 appears in two minimal supports of size 3, whereas e_5 only appears in a minimal support of size 3 and one of size 4; only comparing the cardinals would thus result in a claim that e_0 should necessarily have a higher responsibility than e_1. However, there is another difference between the two in the fact that the completing sets of e_0 intersect whereas those of e_5 don't, and as it turns out the drastic Shapley value gives the score of $\frac{681\,120}{10!}$ to both. In order to compare "all other things being equal" we therefore need to add a hypothesis to the definition of a-domination witness to ensure the intersections between the completing sets are taken into consideration.

A *b-domination witness* $\Delta(\alpha, \beta)$ of α over β is given by:

- an injective mapping $f : \mathrm{CS}(\mathcal{D}, q, \beta) \hookrightarrow \mathrm{CS}(\mathcal{D}, q, \alpha)$;
- for every $S \in \mathrm{CS}(\mathcal{D}, q, \beta)$, an injective mapping $\eta_S : f(S) \hookrightarrow S$;
- a bijection $\eta : \mathcal{D} \setminus \{\alpha\} \hookrightarrow \mathcal{D} \setminus \{\beta\}$ such that for all $S \in \mathrm{CS}(\mathcal{D}, q, \beta)$, and all $\gamma \in f(S)$, $\eta_S(\gamma) = \eta(\gamma)$.

Essentially, the new hypothesis states that the different η_S must agree on the value of any γ. This solves the issue with the example on Figure 6 by preventing the two completing sets of e_5 from mapping to the two completing sets of e_0: say f maps $\{e_6, e_7\}$ to $\{e_1, e_2\}$ and $\{e_8, e_9, e_x\}$ to $\{e_1, e_3\}$; then no correct η would exist because $\eta(e_1)$ would need to be in $\{e_6, e_7\} \cap \{e_8, e_9, e_x\} = \varnothing$.

While this second iteration is a step in the right direction, it isn't quite enough. To see this, recall Example 4.1 from the main body (depicted in Figure 2). In this example, the completing sets for e_0 and e_5 are all disjoint from one another, meaning there exists a b-domination witness of n_0 over n_5, but the completing sets for n_0 intersect some minimal supports (in green) that don't contain it, thereby weakening its contribution (according to the drastic Shapley value at least). We therefore need to add a final condition to avoid this last kind of interference.

A *domination witness* $\Delta(\alpha, \beta)$ of α over β is given by:

- an injective mapping $f : \mathrm{CS}(\mathcal{D}, q, \beta) \hookrightarrow \mathrm{CS}(\mathcal{D}, q, \alpha)$;
- for every $S \in \mathrm{CS}(\mathcal{D}, q, \beta)$, an injective mapping $\eta_S : f(S) \hookrightarrow S$;
- there exists a bijection $\eta : \mathcal{D} \setminus \{\alpha\} \hookrightarrow \mathcal{D} \setminus \{\beta\}$ such that for all S, γ, $\eta_S(\gamma) = \eta(\gamma)$;
- for all $X \in \mathcal{D} \setminus \{\alpha\}$, if X contains some minimal support then $\eta(X)$ does.

(MS1) If there exists a domination witness $\Delta(\alpha, \beta)$ such that one of the $\eta_S : f(S) \hookrightarrow S$ isn't surjective, then $\phi(\mathcal{D}, q, \alpha) > \phi(\mathcal{D}, q, \beta)$.

(MS2) If there exists a domination witness $\Delta(\alpha, \beta)$ such that $f : \mathrm{CS}(\mathcal{D}, q, \beta) \hookrightarrow \mathrm{CS}(\mathcal{D}, q, \alpha)$ isn't surjective, then $\phi(\mathcal{D}, q, \alpha) > \phi(\mathcal{D}, q, \beta)$.

We now show that the drastic Shapley and drastic Banzhaf values satisfiy these axioms indeed.

Proposition B.1. *Let ϕ_c be a drastic Shapley-like score (i.e. a Shapley-like score applied to Ξ^{dr}) with positive coefficient function c. Then ϕ satisfies (Sym-db), (Null-db), (MS1) and (MS2).*

Proof. Recall that drastic Shapley-like scores are defined by the following equation

$$\begin{aligned}\phi_c(\mathcal{D}, q, \alpha) &= \mathcal{S}_c(\mathcal{D}_{\mathsf{n}}, \xi^{\mathrm{dr}}_{q,\mathcal{D}_{\mathsf{x}}}, \alpha)\\ &= \sum_{B \subseteq \mathcal{D}_{\mathsf{n}} \setminus \{\alpha\}} c(|B|, |\mathcal{D}|)(v(B \cup \{\alpha\}) - v(B)) \quad \text{if } \mathcal{D}_{\mathsf{x}} \not\models q, 0 \text{ otherwise} \end{aligned} \tag{4}$$

with c the coefficient function and v the $0/1$ function associated with the query.

(Sym-db) derives from the known axiom (Sym) of Shapley-like scores. For (Null-db): if a player α is irrelevant, then the sum in Equation (4) is empty, hence $\phi_c(\mathcal{D}, q, \alpha) = 0$. Otherwise there is at least a term in the sum; now c has positive values and $v(B \cup \{\alpha\}) - v(B) \in \{0; 1\}$ since q is monotone and Boolean, therefore $\phi_c(\mathcal{D}, q, \alpha) > 0$.

Now assume there exists a domination witness $\Delta(\alpha, \beta)$. We can apply η, which is a bijection, in (4) to express $\phi_c(\mathcal{D}, q, \beta)$ as a sum over subsets of $\mathcal{D} \setminus \{\alpha\}$:

$$\phi_c(\mathcal{D}, q, \beta) = \sum_{X \subseteq \mathcal{D} \setminus \{\alpha\}} c(|\eta(X)|, |\mathcal{D}|)(v(\eta(X) \cup \{\beta\}) - v(\eta(X)))$$

Consider some $X \subseteq \mathcal{D} \setminus \{\alpha\}$ such that $v(\eta(X) \cup \{\beta\}) - v(\eta(X)) \neq 0$, *i.e.* some $\eta(X)$ that contains a completing set S for β and no minimal support. By the definition of domination witnesses, X then contains $\eta^{-1}(S) = \eta_S^{-1}(S) = f(S)$, which is a completing set for α, and no minimal support. In other words, we have $v(X \cup \{\alpha\}) - v(X) \neq 0$, and since the only other possible value is 1 as explained earlier, $v(\eta(X) \cup \{\beta\}) - v(\eta(X)) = v(X \cup \{\alpha\}) - v(X)$, which yields the following:

$$\begin{aligned}\phi_c(\mathcal{D}, q, \beta) &\leq \sum_{X \subseteq \mathcal{D} \setminus \{\alpha\}} c(|X|, |\mathcal{D}|)(v(X \cup \{\alpha\}) - v(X))\\ &= \phi_c(\mathcal{D}, q, \alpha)\end{aligned}$$

At this stage, all it takes for the above inequality to become strict is a coalition X such that $v(\eta(X) \cup \{\beta\}) - v(\eta(X)) < v(X \cup \{\alpha\}) - v(X)$, and this coalition can be obtained from the definition of each axiom.

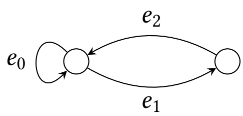

(a)

$$q \mathrel{\hat{=}} \exists xy\, R(x,y) \wedge R(y,x)$$

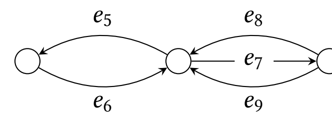

(b)

Fig. 1. Pathological example for some quantitative measures of responsibility. The databases are represented as graphs with nodes being constants and edges being R-tuples.

(MS1) If one of the $\eta_S : f(S) \hookrightarrow S$ isn't surjective, then $\eta(f(S))$ is a proper subset of S, meaning it isn't (nor contains) a completing set for α, therefore $X = f(S)$ gives the desired result.

(MS2) If $f : \mathrm{CS}(\mathcal{D}, q, \beta) \hookrightarrow \mathrm{CS}(\mathcal{D}, q, \alpha)$ isn't surjective, then there is a $X \in \mathrm{CS}(\mathcal{D}, q, \alpha)$ that has no inverse image by f. We can further assume that all η_S are bijective, otherwise the result is already given by (MS1); in other words, for every completing set S for β, $\eta^{-1}(S)$ is a completing set for α. Since η is bijective, if $\eta(X)$ were to contain some $S \in \mathrm{CS}(\mathcal{D}, q, \beta)$, then $\eta^{-1}(S) \subseteq X$ would be a completing set for α, which is contradictory. At last, $\eta(X)$ cannot contain any minimal support because X doesn't. Overall, this X gives the desired result.

□

Note that if one prefers the simpler (and stronger) axioms (MS1a) and (MS2a), which boil down to considering that the interactions between the minimal supports have little impact on responsibility, one can use any WSMS.

Proposition B.2. *If w is positive and strictly decreasing, then ϕ^w_{wsms} satisfies (Sym-db), (Null-db), (MS1a) and (MS2a).*

Proof. (Sym-db) It is a direct consequence of (Sym) and the fact that w is only a function of the minimal supports for q, which are invariant under query equivalence.

(Null-db) If a fact α is irrelevant, then it appears in no minimal support hence the sum in Equation (3) is empty. Otherwise, α appears in some minimal support and the sum is non-empty. Further, all terms will be positive because w is, hence the result.

(MS1a) and (MS2a) If there exists a domination witness $\Delta(\alpha, \beta)$:

$$\begin{array}{rll|l}
\phi^w_{\mathsf{ms}}(\mathcal{D}, q, \beta) & = & \displaystyle\sum_{S \in \mathrm{CS}(\mathcal{D},q,\beta)} w(|S|+1, \mathcal{D}) & \text{by (3)} \\
& \leq & \displaystyle\sum_{S \in \mathrm{CS}(\mathcal{D},q,\beta)} w(|f(S)|+1, \mathcal{D}) & |S| \geq |f(S)| \Rightarrow w(|S|+1,\mathcal{D}) \leq w(|f(S)|+1,\mathcal{D}) \\
& \leq & \displaystyle\sum_{X \in \mathrm{CS}(\mathcal{D},q,\alpha)} w(|X|+1, \mathcal{D}) & \text{because } f \text{ is injective}
\end{array}$$

and the hypothesis of (MS1a) (resp. (MS2a)) makes the second (resp. first) inequality strict. □

B.2 Other missing proofs

It has been claimed in the main text that (MStest) suffices to cover the cases in Figure 1. We formally state and show this claim below.

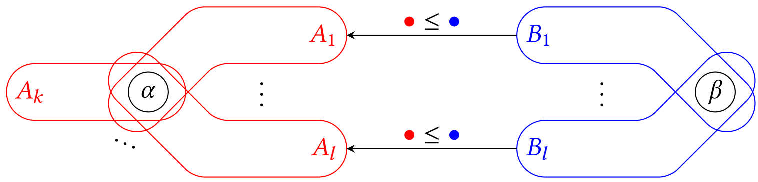

Fig. 3. A visual representation of the setting considered by (MStest). The A_i in red contain α but not β, the B_j in blue contain β but not α, and the minimal supports that contain both have been omitted.

Claim B.3. Consider the example of Figure 1. Then any measure ϕ that satisfies (MStest) will be such that $\phi(\mathcal{D}_{(a)}, q, e_0) > \phi(\mathcal{D}_{(a)}, q, e_1)$ and $\phi(\mathcal{D}_{(b)}, q, e_7) > \phi(\mathcal{D}_{(b)}, q, e_5)$.

Proof. (a) Set $k = l = 1$, $\alpha = e_0$, $\beta = e_1$, $A_1 = \{e_0\}$ and $B_1 = \{e_1, e_2\}$, (MStest) ensures $\phi(\mathcal{D}_{(a)}, q, e_0) > \phi(\mathcal{D}_{(a)}, q, e_1)$.

(b) Set $k = 2$, $l = 1$ $\alpha = e_7$, $\beta = e_5$, $A_1 = \{e_7, e_8\}$, $A_2 = \{e_7, e_9\}$, and $B_1 = \{e_5, e_6\}$. (MStest) ensures $\phi(\mathcal{D}_{(b)}, q, e_7) > \phi(\mathcal{D}_{(b)}, q, e_5)$. □

Proposition 4.2. *If w is positive and strictly decreasing (i.e., such that $k < l \Rightarrow w(k, n) > w(l, n)$ for every n), then ϕ^w_{wsms} satisfies (Sym-db), (Null-db) and (MStest). Further, any drastic Shapley-like score with positive coefficient function satisfies these axioms as well.*[17]

Proof. We check that (MStest) is implied by (the formal version of) (MS1) and (MS2), making this proposition a trivial consequence of Propositions B.1 and B.2.

Take the notations from (MStest), as recalled in Figure 3. The completing sets for α (resp. β) are the $A_i \setminus \{\alpha\}$ (resp. the $B_i \setminus \{\beta\}$), which we shorten to A_i^α (resp. B_i^β)

Define the domination witness $\Delta(\alpha, \beta)$ as follows:

- The injective mapping $f : \mathrm{CS}(\mathcal{D}, q, \beta) \hookrightarrow \mathrm{CS}(\mathcal{D}, q, \alpha)$ is defined by $f(B_i^\beta) = A_i^\alpha$ for every $i \leq l$.
- For every $i \leq l$, $|A_i| \leq |B_i|$ implies the existence of some injective mapping $\eta_{B_i^\beta} : A_i^\alpha \hookrightarrow B_i^\beta$.
- The only elements that are changed by some η_S are the elements of some A_i^α, which only appear in one of them, hence only have one possible image. This means that the bijection $\eta : \mathcal{D} \setminus \{\alpha\} \hookrightarrow \mathcal{D} \setminus \{\beta\}$ that swaps α with β, every $\gamma \in A_i^\alpha$ with $\eta_{A_i^\alpha}(\gamma)$ and leaves the rest unchanged is such that for all S, γ, $\eta_S(\gamma) = \eta(\gamma)$.
- For all $X \in \mathcal{D} \setminus \{\alpha\}$, if X contains some minimal support then it can only be some B_i, because all others contain α. Now by definition η swaps every element of A_i with some element of B_i meaning that $\eta(B_i) \subseteq \eta(X)$ will contain the minimal support A_i.

As shown above, $\Delta(\alpha, \beta)$ is a domination witness indeed. We conclude with the last hypothesis of (MStest): if $\exists i.|A_i| < |B_i|$ then $\eta_{A_i^\alpha}$ isn't surjective and (MS1) gives the desired result, and if $k > l$ then f isn't surjective and (MS2) does. □

Proposition 4.4. *For every WSMS ϕ^w_{wsms} and every Shapley-like score $\mathcal{S}_c$ whose coefficient function c is positive on all inputs, there exists a family $\Xi \triangleq \left(\xi_{q,\mathcal{D}_x}\right)$ such that for every query q, database $\mathcal{D}$ and $\alpha \in \mathcal{D}$, $\phi^w_{\text{wsms}}(\mathcal{D}, q, \alpha) = \mathcal{S}_c(\mathcal{D}_n, \xi_{q,\mathcal{D}_x}, \alpha)$.*

[17] All these measures also satisfy (the formal versions of) (MS1) and (MS2) by Propositions B.1 and B.2 in the appendix.

Proof. To obtain a WSMS, we shall use a score function $\xi^{\gamma}_{\mathsf{MS}_q(\mathcal{D}),\mathcal{D}_\mathsf{x}}$ which itself is a weighted sum of minimal supports in some sense. Note that to simplify the notations later we directly take the set of minimal supports as a parameter instead of the query.

$$\begin{aligned}\xi^{\gamma}_{\mathsf{MS}_q(\mathcal{D})} &: 2^{\mathcal{D}_\mathsf{n}} \longrightarrow \mathbb{R}\\ &B \mapsto \sum_{\substack{S\in\mathsf{MS}_q(\mathcal{D})\\ S\subseteq B}} \gamma(S,\mathcal{D}_\mathsf{n},\mathcal{D}_\mathsf{x})\end{aligned}$$

The coefficients $\gamma(S,\mathcal{D}_\mathsf{n},\mathcal{D}_\mathsf{x})$ are to be set later.

It is easy to see that any Shapley-like score follows the linearity axiom $\mathcal{S}_c(\mathcal{D}_\mathsf{n},\xi,\alpha)+\mathcal{S}_c(\mathcal{D}_\mathsf{n},\xi',\alpha) = \mathcal{S}_c(\mathcal{D}_\mathsf{n},\xi+\xi',\alpha)$. From this we can get:

$$\mathcal{S}_c(\mathcal{D}_\mathsf{n},\xi^{\gamma}_{\mathsf{MS}_q(\mathcal{D})},\alpha) = \sum_{S\in\mathsf{MS}_q(\mathcal{D})} \mathcal{S}_c(\mathcal{D}_\mathsf{n},\xi^{\gamma}_{\{S\}},\alpha) \tag{5}$$

We now fix some $S \in \mathsf{MS}_q(\mathcal{D})$, and denote by S_n its endogenous part. If $\alpha \notin S_\mathsf{n}$, it is easy to see that $\mathcal{S}_c(\mathcal{D}_\mathsf{n},\xi^{\gamma}_{\{S\}},\alpha) = 0$. Now consider $\alpha \in S_\mathsf{n}$, and denote $S^-_\mathsf{n} \mathrel{\hat{=}} S_\mathsf{n} \setminus \{\alpha\}$. By applying (2) we get

$$\begin{aligned}\mathcal{S}_c(\mathcal{D}_\mathsf{n},\xi^{\gamma}_{\{S\}},\alpha) &= \sum_{\substack{B\subseteq\mathcal{D}_\mathsf{n}\setminus\{\alpha\}\\ S^-_\mathsf{n}\subseteq B}} c(|B|,|\mathcal{D}_\mathsf{n}|)\gamma(S,\mathcal{D}_\mathsf{n},\mathcal{D}_\mathsf{x})\\ &= \sum_{k=0}^{|\mathcal{D}_\mathsf{n}|-1} c(k,|\mathcal{D}_\mathsf{n}|)\gamma(S,\mathcal{D}_\mathsf{n},\mathcal{D}_\mathsf{x})|\{B\subseteq\mathcal{D}_\mathsf{n}\setminus\{\alpha\} \mid S^-_\mathsf{n}\subseteq B \wedge |B|=k\}|\end{aligned}$$

Computing the rightmost factor boils down to choosing the $k-|S^-_\mathsf{n}|$ elements of $B\setminus S_\mathsf{n}$ out of the $|\mathcal{D}_\mathsf{n}|-|S_\mathsf{n}|$ possible ones, since there is no choice on the others. This gives

$$\mathcal{S}_c(\mathcal{D}_\mathsf{n},\xi^{\gamma}_{\{S\}},\alpha) = \gamma(S,\mathcal{D}_\mathsf{n},\mathcal{D}_\mathsf{x}) \sum_{k=|S_\mathsf{n}|-1}^{|\mathcal{D}_\mathsf{n}|-1} \binom{|\mathcal{D}_\mathsf{n}|-|S_\mathsf{n}|}{k-|S_\mathsf{n}|+1} c(k,|\mathcal{D}_\mathsf{n}|)$$

Hence by setting

$$\gamma(S,\mathcal{D}_\mathsf{n},\mathcal{D}_\mathsf{x}) \mathrel{\hat{=}} \frac{w(|S|,|\mathcal{D}_\mathsf{n}|)}{\sum_{k=|S_\mathsf{n}|-1}^{|\mathcal{D}_\mathsf{n}|-1}\binom{|\mathcal{D}_\mathsf{n}|-|S_\mathsf{n}|}{k-|S_\mathsf{n}|+1}c(k,|\mathcal{D}_\mathsf{n}|)}$$

we obtain $\mathcal{S}_c(\mathcal{D}_\mathsf{n},\xi^{\gamma}_{\mathsf{MS}_q(\mathcal{D})},\alpha) = \phi^{w}_{\mathsf{wsms}}(\mathcal{D},q,\alpha)$ by Equation (5), as desired. Note that there cannot be a division by 0 in the definition of γ because the coefficient function c is assumed to be positive on all inputs. □