Appendix A Proof Details

Proof of Lemma 3. Let G be a ground atom with $\nu(G) > 0$. Since ν is minimal there must be a ground rule with G in the head such that for every body atom $G_i \in \mathrm{body}(\gamma)$ also $\nu(G_i) > 0$. Consider the tree with γ as the root and a child for every body atom of γ. At each child take again a rule where the respective body atom is in the head and the body is true to some degree greater than 0.

Repeating the process, we must arrive at a finite tree where all leaves must correspond to ground rules where all body atoms are defined in τ with truth greater than 0. Otherwise, we could simply set $\nu(G') = 0$ for all ground atoms G' whose branches do not lead to such leaves and ν would still be a K-fuzzy model, contradicting its minimality.

Viewing the same tree in the context of Π^{crisp}, D_τ we then see, going from the leaves upward that the bodies of all these rules must be true in all models. Thus, also ultimately $\Pi^{crisp}, D_\tau \models G$. ☐

Proof of Lemma 5. Let $\mathbf{x}$ be a feasible solution of $Opt^K_{\Pi,\tau}$. Observe that $\mathrm{Eval}(\gamma) = \nu_{\mathbf{x}}(\gamma)$ for every ground rule γ: Since, $\mathrm{Eval}(\gamma) \geq K$ for all $\gamma \in \Gamma$, we also see that $\nu_{\mathbf{x}}$ K-satisfies every rule in $OGround(\Pi^{crisp}, D_\tau)$. By definition of an oblivious chase sequence and since $\nu_{\mathbf{x}}(G) = 0$ for any ground atom that is not mentioned in $OLim(\Pi^{crisp}, D_\tau)$, any other ground rule $\gamma \notin \Gamma$ is trivially K-satisfied since both head and body have truth 0.

Let ν be a K-fuzzy model of (Π, τ). Let $\mathbf{x}$ be a solution of $Opt^K_{\Pi,\tau}$ with $x_i = \nu(G_i)$ for all $G \in \mathcal{G}$. Since ν K-satisfies all ground rules, we in particular have that $\nu(\gamma) = \mathrm{Eval}(\gamma) \geq K$ for all $\gamma \in \Gamma$ and therefore $\mathbf{x}$ is feasible. The third statement follows immediately from combination of the first two. ☐

Proof of Lemma 6. By Lemma 3, and the fact that the oblivious chase sequence constructs minimal models for Datalog, we also have that $\mathcal{G}' \subseteq OLim(\Pi^{crisp}, D_\tau)$. It follows by construction that every derived $G \in \mathcal{G}$ in the head of some ground rule in $OGround(\Pi^{crisp}, D_\tau)$. What is left to show is that at least one such rule is also ν-tight.

Suppose the statement is false and let $\delta > 0$ be the minimal ν-gap of rules in Γ whose body atoms are all not in $\mathcal{G}'$. Let Γ' contain only the ν-tight rules in $OGround(\Pi^{crisp}, D_\tau)$ that contain an atom from $\mathcal{G}'$ in the head. Since we assume the statement false, all of the rules in Γ' have a body atom that is also in $\mathcal{G}'$. Now consider the truth assignment ν' defined as

$$\nu'(G) = \begin{cases} \max\{0, \nu(G) - \delta\} & \text{if } G \in \mathcal{G}' \\ \nu_{\mathbf{x}}(G) & \text{otherwise} \end{cases}$$

By assumption we have that all ν-tight rules where $\mathcal{G}'$ occurs in the head also have some atom from $\mathcal{G}'$ in the body. Meaning that under ν' the truth of both head and body are decreased by at least δ (or until both are 0). Hence, all rules in Γ' are still K-satisfied. We can also see that all other rules that contain an atom in $\mathcal{G}'$ remain K-satisfied under ν'. When an atom of $\mathcal{G}'$ occurs in the body of a rule, the implication can only become more true. Where an atom G from $\mathcal{G}'$ occurs only in the head of a rule $\gamma \notin \Gamma'$, we have that the $\nu_{\mathbf{x}}$-gap of the rule is at least δ. That is,

$$\nu_{\mathbf{x}}(G) \geq \nu_{\mathbf{x}}(\mathrm{body}(\gamma)) - 1 + K + \delta$$

and therefore also

$$\nu'(G) \geq \nu_{\mathbf{x}}(G) - \delta \geq \nu_{\mathbf{x}}(\text{body}(\gamma)) - 1 + K \geq \nu'(\text{body}(\gamma)) - 1 + K$$

We therefore see that ν' is a K-fuzzy model with $\nu' < \nu$ and we arrive at a contradiction. □

Proof of Theorem 8. The implication from right to is a a special case of Lemma 3. For the other direction, we argue by induction over the oblivious chase sequence $D_0, D_1, \ldots$ for Π^{crisp}, D^1, that if some ground atom G is in D_i, then $\mu(G) = 1$ in every 1-fuzzy model μ of (Π, τ). The base case $D_0 = D^1$ follows by definition of D^1. For the induction, suppose the claim holds up to step i. In the step from D_i to D_{i+1} either the ground atoms in both are the same (the oblivious application produced a ground head that was already in D_i) or there is a single new ground atom G in D_{i+1} but not in D_i. Let γ be the ground rule induced by the oblivious application in the step from i to $i+1$. Since the rule was applicable, for all $G' \in \text{body}(\gamma)$ we have $G' \in D_i$ and by the induction hypothesis also $\mu(G') = 1$. Hence, $\mu(\text{body}(\gamma)) = 1$ and since γ must be 1-satisfied, we have that in every 1-fuzzy model also $\mu(G) = 1$.

Recall, that $OLim(\Pi^{crisp}, D^1)$ is always a minimal model of Π^{crisp}, D^1 and thus $\Pi^{crisp}, D^1 \models G$ if and only if $G \in OLim(\Pi^{crisp}, D^1)$. And by the above induction argument $G \in OLim(\Pi^{crisp}, D^1)$ implies that $\mu(G) = 1$. □

Proof of Theorem 9. First, observe that since $OLim(\Pi^{crisp}, D_\tau)$ is unique up to isomorphism, we can assume without loss of generality that there is no preferred model where some ground atom G is true but not in the set $\mathcal{G}$ of ground atoms considered in the construction of $\exists\text{-}Opt^K_{\Pi,\tau}$.

Let ν be a preferred K-fuzzy model of Π, τ. Since ν has an oblivious base, we have that $\nu(G) > 0$ only if $G \in \mathcal{G}$. Let $\mathbf{x}$ be the solution of $\exists\text{-}Opt^K_{\Pi,\tau}$ where $x_i = \nu(G_i)$ for all $G_i \in \mathcal{G}$. Then, by definition of Eval we immediately see $\text{Eval}(\gamma) = \nu(\gamma)$ for every ground rule γ. Hence, if all ground rules are K-satisfied by ν, then all constraints $\text{Eval}(\gamma) \geq K$ are satisfied and ν is feasible.

For the second statement, assume that $\mathbf{x}$ is an optimal solution of $\exists\text{-}Opt^K_{\Pi,\tau}$ but $\nu_{\mathbf{x}}$ is not preferred. The only way $\nu_{\mathbf{x}}$ can not be preferred is if it is not active minimal, i.e., there is some K-fuzzy model μ where for all $G \in GAtoms[Adom]$, $\mu(G) \leq \nu_{\mathbf{x}}(G)$ and for at least one $G' \in GAtoms[Adom]$, $\mu(G') < \nu_{\mathbf{x}}(G')$. Since both μ and $\nu_{\mathbf{x}}$ have an oblivious base (which is unique up to isomorphism) we see that every $G \in GAtoms$ where $\mu(G) > 0$ is also in $\mathcal{G}$. Hence, it is straightforward to construct a feasible solution $\mathbf{x}'$ from μ for which the objective function is strictly lower than for $\mathbf{x}$, a contradiction. □

Appendix B Finiteness in the Oblivious Chase

Since we are interested particularly in instances where the chase is finite, it is of interest to identify fragments where this is always the case. The most prominent condition for which the finiteness of the chase in Datalog$^{\pm}$ is studied, is the restriction to *weakly acyclic programs* as first introduced by Fagin et al. (2005). Let Π be a set of rules with over the relational language σ. We first define the *dependency graph* for Π as the graph with vertices $\{(R,i) \mid R \in \sigma, i \in \{1,\ldots,\#(R)\}\}$ We say a variable x is in position (R,i) if x is at the i-th place of a relation with name R. There is a (normal) edge from vertex (R,i) to

(S, j) exactly if there is a rule in Π where the same variable x occurs at position (R, i) in the body and at position (S, j) in the head of the rule. There is a *special edge* from (R, i) to (S, j) exactly if there is an existential rule in Π, such that there is a variable x that x occurs in both the body and head, x is in position (R, i) of the body, and an existentially quantified y is in position (S, j) in the head. We say that a program is *weakly acyclic* if its dependency graph has no cycle that passes a special edge.

However, the standard definition of weakly acyclic is particular to the restricted chase, which does not fit our setting. This becomes apparent when considering the simple (weakly acyclic) program

$$P(x) \to \exists y\ P(y).$$

The oblivious chase is infinite in this case, since every new instantiation of y generates a new homomorphism that satisfies the body. Since x does not occur in the head at all, the dependency graph used in the standard definition of weakly-acyclic does not have any edges and is thus weakly acyclic.

Let the *variable expansion* $\text{ve}(\Pi)$ of Π be the program obtained by the following rewriting. Without loss of generality every existential rule is of the form $\varphi(\mathbf{x}) \to \exists \mathbf{y}\ R(\mathbf{y}, \mathbf{x}')$, where φ is the body formula, R is a relation symbol and $\mathbf{x}'$ are those (free) variables that occur in both the body and the head. Let $\mathbf{x}^*$ be the free variables that only occur in φ but not in the head. Then, for every such rule we replace it by the two rules

$$\begin{aligned} &\varphi(\mathbf{x}) && \to \exists \mathbf{y}\ R^*(\mathbf{y}, \mathbf{x}', \mathbf{x}^*) \\ &R^*(\mathbf{y}, \mathbf{x}', \mathbf{x}^*) && \to R(\mathbf{y}, \mathbf{x}') \end{aligned}$$

to obtain $\text{ve}(\Pi)$. Intuitively, the variable expansion reveals those particular cycles that are harmless in the restricted chase but lead to infinite sequences in the oblivious chase in the standard definition of weak acyclicity. For example, the variable expansion of our simple example above would thus be the (no longer weakly acyclic) program

$$\begin{aligned} &P(x) && \to \exists y\ P^*(y, x) \\ &P^*(y, x) && \to P(y) \end{aligned}$$

The following theorem then follows by similar argument as originally given for the restricted chase (Fagin et al. 2005, Theorem 3.9) by additionally observing that nulls can never propagate to a position in the dependency graph with lower rank, where the rank of a position is the maximal number of special edges on an incoming path. One can then inductively bound the number of possible groundings of the body to consequently bound the number of nulls generated in the head for each position.

Theorem 10. *Let Π, τ be a MV$^\pm$-instance program. If $\text{ve}(\Pi)$ is weakly acyclic, then $OLim(\Pi^{crisp}, D_\tau)$ is finite and of polynomial size in data complexity.*