q that mutually attack one another. The proof is a first-order reduction from CERTAINTY(q_0) with $q_0 = \{R_0(\underline{x}, y), S_0(\underline{y}, x)\}$, which is an **L**-hard problem [17].

For every pair (a, b) of constants, define Θ^a_b as the valuation over $\mathsf{vars}(q)$ such that for every $x \in \mathsf{vars}(q)$,

$$\Theta^a_b(x) = \begin{cases} a & \text{if } x \in [R:K]^{\ominus,q} \setminus [S:L]^{\ominus,q} \\ b & \text{if } x \in [S:L]^{\ominus,q} \setminus [R:K]^{\ominus,q} \\ \bot & \text{if } x \in [R:K]^{\ominus,q} \cap [S:L]^{\ominus,q} \\ \langle a, b\rangle & \text{otherwise} \end{cases}$$

Claim D.1. *$\{\Theta^a_b(R), \Theta^{a'}_{b'}(R)\}$ is inconsistent if and only if $a = a'$ and $b \neq b'$.*

Proof of Claim D.1. $\boxed{\Longrightarrow}$ By contraposition, it suffices to show the following:

(A) if $a \neq a'$, then $\{\Theta^a_b(R), \Theta^{a'}_{b'}(R)\}$ is consistent; and
(B) if $b = b'$, then $\{\Theta^a_b(R), \Theta^{a'}_{b'}(R)\}$ is consistent.

For (A), assume $a \neq a'$. We have that $K \neq \emptyset$, or else $[\![R : K]\!]$ would have no incoming attack, a contradiction. Since $K \subseteq [R : K]^{\ominus,q}$, it is obvious that $\{\Theta^a_b(R), \Theta^{a'}_{b'}(R)\} \models [\![R : K]\!]$.

Let $[\![R : M]\!]$ be another key of q. By our definition of Θ^a_b, it suffices to show $M \nsubseteq [S : L]^{\ominus,q}$. Assume, toward a contradiction, $M \subseteq [S : L]^{\ominus,q}$, that is, $\mathcal{K}(q \ominus S) \models L \rightarrow M$. Since $\mathcal{K}(q \ominus S) \models M \rightarrow K$ is obvious, we have $\mathcal{K}(q \ominus S) \models L \rightarrow K$. But then $[\![S : L]\!]$ cannot attack $[\![R : K]\!]$, a contradiction.

For (B), assume $b = b'$. If $a \neq a'$, then the desired result follows from (A). If $a = a'$, then $\Theta^a_b(R) = \Theta^{a'}_{b'}(R)$, and the desired result follows vacuously (because a database instance that is a singleton satisfies every KD).

$\boxed{\Longleftarrow}$ Assume $a = a'$ and $b \neq b'$. From $K \subseteq [R : K]^{\ominus,q}$, it follows that Θ^a_b and $\Theta^{a'}_{b'}$ agree on K. Since $[\![R : K]\!]$ has an outgoing attack, $\mathsf{vars}(R) \nsubseteq [R : K]^{\ominus,q}$, and the desired result obtains because Θ^a_b and $\Theta^{a'}_{b'}$ disagree on variables in $\mathsf{vars}(R) \setminus [R : K]^{\ominus,q}$. ∎

Claim D.2. *If $\Theta^a_b(R) = \Theta^{a'}_{b'}(R)$, then $a = a'$ and $b = b'$.*

Proof of Claim D.2. Assume $\Theta^a_b(R) = \Theta^{a'}_{b'}(R)$. As argued in the proof of Claim D.1, $\emptyset \neq K \subseteq [R : K]^{\ominus,q}$. It obviously follows $a = a'$.

Since $\{\Theta^a_b(R), \Theta^{a'}_{b'}(R)\}$ is consistent, Claim D.1 implies $a \neq a'$ or $b = b'$. Since $a = a'$, if follows $b = b'$. ∎

Claim D.3. *For every atom $T \in q \setminus \{R, S\}$, for all constants a, a', b, b', $\{\Theta^a_b(T), \Theta^{a'}_{b'}(T)\}$ is consistent.*

Proof of Claim D.3. Let $[\![T : M]\!]$ be a key of q. Assume that Θ^a_b and $\Theta^{a'}_{b'}$ agree on M. We need to show $\Theta^a_b(T) = \Theta^{a'}_{b'}(T)$. We have that both $\mathcal{K}(q \ominus R)$ and $\mathcal{K}(q \ominus S)$ contain $M \rightarrow \mathsf{vars}(T)$. The desired result is obvious if $a = a'$ and $b = b'$. Three other cases can occur:

- Case $a = a'$ and $b \neq b'$. Then $M \subseteq [R : K]^{\ominus,q}$, hence $\mathcal{K}(q \ominus R) \models K \rightarrow M$. It follows $\mathcal{K}(q \ominus R) \models K \rightarrow \mathsf{vars}(T)$, and therefore $\mathsf{vars}(T) \subseteq [R : K]^{\ominus,q}$. The desired result follows from the definition of Θ^a_b.
- Case $a \neq a'$ and $b = b'$. This case is symmetrical to the previous item.
- Case $a \neq a'$ and $b \neq b'$. Then $M \subseteq [R : K]^{\ominus,q} \cap [S : L]^{\ominus,q}$, hence $\mathcal{K}(q \ominus R) \models K \rightarrow M$ and $\mathcal{K}(q \ominus S) \models L \rightarrow M$. It follows $\mathcal{K}(q \ominus R) \models K \rightarrow \mathsf{vars}(T)$ and $\mathcal{K}(q \ominus S) \models L \rightarrow \mathsf{vars}(T)$, and therefore $\mathsf{vars}(T) \subseteq [R : K]^{\ominus,q} \cap [S : L]^{\ominus,q}$. The desired result follows from the definition of Θ^a_b.

This concludes the proof of Claim D.3. ∎

For every instance **db** of CERTAINTY(q_0), we define $f(\mathbf{db})$ as the following database instance:

- for every $R_0(\underline{a}, b) \in \mathbf{db}$ such that $S_0(\underline{b}, a) \notin \mathbf{db}$, $f(\mathbf{db})$ includes $\Theta^a_b(q \setminus \{S\})$;
- for every $R_0(\underline{a}, b) \in \mathbf{db}$ such that $S_0(\underline{b}, a) \in \mathbf{db}$, $f(\mathbf{db})$ includes $\Theta^a_b(q)$; and
- for every $S_0(\underline{b}, a) \in \mathbf{db}$ such that $R_0(\underline{a}, b) \notin \mathbf{db}$, $f(\mathbf{db})$ includes $\Theta^a_b(q \setminus \{R\})$.

Let $F(\mathbf{db})$ be the subset of $f(\mathbf{db})$ that contains all facts that are neither R-facts nor S-facts. By Claims D.1, D.2, and D.3 (and their symmetric versions),

$$\mathsf{rset}(f(\mathbf{db})) = \{f(\mathbf{r}) \cup F(\mathbf{db}) \mid \mathbf{r} \in \mathsf{rset}(\mathbf{db})\}.$$

Let **db** be an instance of CERTAINTY(q_0). It suffices to show that for every repair **r** of **db**,

$$\mathbf{r} \models q_0 \text{ if and only if } f(\mathbf{r}) \cup F(\mathbf{db}) \models q.$$

To show the latter equivalence, let **r** be an arbitrary repair of **db**.

$\boxed{\Longrightarrow}$ Assume $\mathbf{r} \models q_0$. There exist constants a, b such that $R_0(\underline{a}, b), S_0(\underline{b}, a) \in \mathbf{r}$. It follows that $f(\mathbf{r})$ includes $\Theta^a_b(q)$, hence $f(\mathbf{r}) \cup F(\mathbf{db}) \models q$.

$\boxed{\Longleftarrow}$ Assume $f(\mathbf{r}) \cup F(\mathbf{db}) \models q$. There is a valuation μ such that $\mu(q) \subseteq f(\mathbf{r}) \cup F(\mathbf{db})$. Then, there are $R_0(\underline{a}, b), S_0(\underline{b'}, a') \in \mathbf{r}$ such that $\mu(R) = \Theta^a_b(R)$ and $\mu(S) = \Theta^{a'}_{b'}(S)$. It suffices to show $a = a'$ and $b = b'$. We show $a = a'$; the proof of $b = b'$ is symmetrical. Since $[\![R : K]\!] \overset{q}{\rightsquigarrow} [\![S : L]\!]$, there is a sequence

$$F_0, v_1, F_1, v_2, F_2 \dots, F_i, v_{i+1}, F_{i+1}, \dots, v_\ell, F_\ell$$

where $\{F_i\}_{i=0}^{\ell} \subseteq q$, $F_0 = R$, $F_\ell = S$, and for every $i \in \{1, \dots, \ell\}$, $v_i \in \mathsf{vars}(q) \setminus [R : K]^{\ominus,q}$ such that v_i occurs in both F_{i-1} and F_i. For every $i \in \{0, \dots, \ell\}$, there are constants a_i, b_i such that $\mu(F_i) = \Theta^{a_i}_{b_i}(F_i)$. We show by induction on increasing i that for every $i \in \{0, \dots, \ell\}$, $b_i = b$.

Basis $i = 0$. Since $\Theta^a_b(R) = \mu(R) = \Theta^{a_0}_{b_0}(R)$, it follows by Claim D.2 that $b_0 = b$.

Step $i \rightarrow i + 1$. The induction hypothesis is $b_i = b$. We have that μ, $\Theta^{a_i}_{b_i}$, and $\Theta^{a_{i+1}}_{b_{i+1}}$ agree on v_{i+1}. Then, from $\Theta^{a_i}_{b_i}(v_{i+1}) \in \{b, \langle a_i, b\rangle\}$ and $\Theta^{a_{i+1}}_{b_{i+1}}(v_{i+1}) \in \{b_{i+1}, \langle a_{i+1}, b_{i+1}\rangle\}$, it follows $b = b_{i+1}$.

Consequently, $b_\ell = b$. Since $\Theta^{a'}_{b'}(S) = \mu(S) = \Theta^{a_\ell}_{b}(S)$, it follows from (the symmetrical of) Claim D.2 that $b = b'$. This concludes the proof of Lemma 5.5. □