Theorem 0.1. *For every constant integer $k \geq 2$ and constant $\eta > 0$, there exists a large enough constant $q = q(k, \eta)$ such that the following problem is* NP*-hard: Given a Consistent Labeling instance $\mathcal{H}$ with alphabet size q, distinguish between* $\mathrm{Opt}^{\mathrm{strong}}(\mathcal{H}) = 1$ *and* $\mathrm{Opt}^{\mathrm{weak}}(\mathcal{H}) < \eta$.

Proof. Given k and η, let q be a large enough constant so that Label Cover with alphabet size q has NP-hardness of 1 vs. $\eta/\binom{k}{2}$. We now reduce from this problem. Given a Label Cover instance $\mathcal{G}$ with left vertices U and right vertices V, our Consistent Labeling instance will have vertex set V. If the Label Cover instance has label set Σ, we take $\mathcal{H}$'s key and label sets to be $K = L = \Sigma$. For each ordered k-tuple of Label Cover constraints $[(u, w_1), \ldots, (u, w_k)]$ emanating from the same vertex $u \in U$ we produce a hyperedge constraint $e = (w_1, \ldots, w_k)$ in $\mathcal{H}$. (This indeed gives a regular hypergraph.) The associated maps $\pi_e^i : L \to K$ are just the maps $\pi_{w_i \to u}$ from $\mathcal{G}$. This fully specifies the Consistent Labeling instance $\mathcal{H}$. The reduction is clearly polynomial time.

Completeness is trivial; if $f : U \cup V \to \Sigma$ is an assignment satisfying all of $\mathcal{G}$'s constraints, f's restriction to V is a labeling strongly satisfying all of $\mathcal{H}$'s constraints.

As for soundness, suppose $f : V \to L$ is a labeling weakly satisfying at least an η fraction of the hyperedge constraints. This implies that if we pick a random $u \in U$ and k of its neighbors $v_1, \ldots, v_k$ independently, then with probability at least η, at least two of $\pi_{v_i \to u}(f(v_i))$ are the same. It follows that if we pick a random u and then two independent neighbors v_1 and v_2, then with probability at least $\eta/\binom{k}{2}$ we have $\pi_{v_1 \to u}(f(v_1)) = \pi_{v_1 \to u}(f(v_1))$. (Think of picking k neighbors first, then choosing two of them further at random.) Equivalently, if we pick a random (u, v) in G, and then pick another random neighbor v' of u, the probability that $\pi_{v \to u}(f(v)) = \pi_{v' \to u}(f(v'))$ is at least $\eta/\binom{k}{2}$.

Now let us randomly extend f to a map $f : U \cup V \to \Sigma$ by setting each $f(u)$ to $\pi_{v' \to u}(f(v'))$ where v' is a random neighbor of u. Then we have equivalently shown that if (u, v) is a random constraint in $\mathcal{G}$ and f is chosen randomly, the probability that f satisfies (u, v) is at least $\eta/\binom{k}{2}$. I.e, the expected fraction of constraints f satisfies is at least $\eta/\binom{k}{2}$. Hence there exists an $f : U \cup V \to \Sigma$ satisfying at least this fraction of constraints of $\mathcal{G}$, and hence $\mathrm{Opt}(\mathcal{G}) \geq \eta/\binom{k}{2}$. $\square$