A. Extra Proofs

Proof of Proposition 3. It is well known that the local sensitivity of any contingency table with respect to our definition of $\text{nbrs}(\mathbf{X})$ is one. This is easy to see from the definition of $\mathbf{n}_C$ following Eq. (2): each individual contributes a count of exactly one to each clique contingency table. Since there are $|\mathcal{C}|$ tables, the local sensitivity is exactly $|\mathcal{C}|$ for all data sets, and, therefore, the sensitivity is the same. □

Proof of Proposition 4. Note that $n_C(i_C)$ is a sum of N iid indicator variables, so $n_C(i_C) \sim \text{Binomial}\big(N, \mu_C(i_C)\big)$, and $\text{Var}\left(n_c(i_C)\right) = N\mu_C(i_C)\big(1 - \mu_C(i_C)\big)$. Now let $z \sim \text{Laplace}(|\mathcal{C}|/\epsilon)$ and write:

$$\bar{\mu}_C(i_C) = \frac{1}{N}\big(n_C(i_C) + z\big)$$

Recall that $\mathbb{E}[z] = 0$ and $\text{Var}(z) = 2|\mathcal{C}|^2/\epsilon^2$. We see immediately that $\mathbb{E}[\bar{\mu}_C(i_C)] = \mathbb{E}\big[n_C(i_C)/N\big] = \mu_C(i_C)$. Therefore, the estimator is unbiased and its mean-squared error is equal to its variance. Since $n_C(i_C)$ and z are independent, we have:

$$\begin{aligned}\text{Var}\left(\bar{\mu}_C(i_C)\right) &= \frac{\text{Var}\left(n_C(i_C)\right)}{N^2} + \frac{\text{Var}(z)}{N^2} \\ &= \frac{\mu_C(i_C)\big(1 - \mu_C(i_C)\big)}{N} + \frac{2|\mathcal{C}|^2}{N^2\epsilon^2}\end{aligned}$$

The fact that $p(\mathbf{x}; \hat{\boldsymbol{\theta}})$ converges to $p(\mathbf{x}; \boldsymbol{\theta})$ follows from Proposition 2 and the consistency of the marginals, as long as the true marginals $\boldsymbol{\mu}$ lie in the interior of the marginal polytope $\mathcal{M}$. However, this is guaranteed because the true distribution $p(\mathbf{x}; \boldsymbol{\theta})$ is strictly positive. □

Proof of Proposition 5. After applying Stirling's approximation to $\log p(\mathbf{n}; \boldsymbol{\theta})$ we obtain (Nguyen et al., 2016):

$$\log h(\mathbf{n}) \approx H(\mathbf{n}) = N \log N + \sum_{C \in \mathcal{C}} \hat{H}_C - \sum_{S \in \mathcal{S}} \nu(S)\hat{H}_S \tag{7}$$

where we define $\hat{H}_A = -\sum_{i_A \in \mathcal{X}^{|A|}} n_A(i_A) \log n_A(i_A)$ for any $A \in \mathcal{C} \cup \mathcal{S}$. The term $\hat{H}_A$ is a scaled entropy. We can rewrite it as:

$$\begin{aligned}\hat{H}_A &= -N \sum_{i_A} \frac{n_A(i_A)}{N} \log\Big(\frac{n_A(i_A)}{N} \cdot N\Big) \\ &= -N \sum_{i_A} \hat{\mu}_A(i_A) \log \hat{\mu}_A(i_A) - N \sum_{i_A} \hat{\mu}_A(i_A) \log N \\ &= N H_A - N \log N\end{aligned}$$

where H_A is now the entropy of the empirical marginal distribution $\hat{\boldsymbol{\mu}}_A = \mathbf{n}_A/N$. Since the total multiplicity of the separators is one less than the number of cliques, when we substitute back into Eq. (7), all of the $N \log N$ terms cancel, and we are left only with

$$H(\mathbf{n}) = N \cdot \Big(\sum_{C \in \mathcal{C}(\mathcal{T})} H_A - \sum_{S \in \mathcal{S}(\mathcal{T})} \nu(S) H_A\Big)$$

But, from standard arguments about the decomposition of entropy on junction trees, the term in parentheses is exactly the entropy of distribution q defined as:

$$q(\mathbf{x}) = \frac{\prod_{C \in \mathcal{C}} \prod_{i_C \in \mathcal{X}^{|C|}} \hat{\mu}_C(\mathbf{x}_C)}{\prod_{S \in \mathcal{S}} \prod_{i_S \in \mathcal{X}^{|S|}} \hat{\mu}_S(\mathbf{x}_S)^{\nu(S)}},$$

which factors according to $\mathcal{C}$ and can be written as $p(\mathbf{x}; \boldsymbol{\theta})$ for parameters $\boldsymbol{\theta}$ derived from the marginal probabilities. Although the mapping from parameters to distributions is many-to-one, for any maginals $\hat{\boldsymbol{\mu}}$, there is a unique distribution $p(\mathbf{x}; \boldsymbol{\theta})$ in the model family that has marginals $\hat{\boldsymbol{\mu}}$ (Wainwright & Jordan, 2008), so this uniquely defines $q(\mathbf{x})$ as stated in the Proposition.

□