Proof that $|\mathcal{K}| \geq |\mathcal{M}|$ is necessary for perfect privacy

Let's first remind ourselves of the definition of perfect privacy. Here we use the alternative, but equivalent, formulation given in Problem 4 of Problem set 1.

Definition 1. An encryption scheme $\Pi = (\mathsf{Enc}, \mathsf{Dec})$ defined over $(\mathcal{K}, \mathcal{M}, \mathcal{C})$ has *one-time perfect privacy* if for every $M, M' \in \mathcal{M}$ and every $C \in \mathcal{C}$

$$\Pr[\mathsf{Enc}_K(M) = C] = \Pr[\mathsf{Enc}_K(M') = C].$$

In both cases the probability is taken over the random choice $K \xleftarrow{\$} \mathcal{K}$ and the coins tossed by Enc (if any).

Theorem 1. *No symmetric encryption scheme has perfect privacy if $|\mathcal{K}| < |\mathcal{M}|$.*

Proof. Let $\Pi = (\mathsf{Enc}, \mathsf{Dec})$ be a symmetric encryption scheme and assume $|\mathcal{K}| < |\mathcal{M}|$. We want to show that Π cannot have perfect privacy.

By the definition of perfect privacy, in order to prove the theorem it is sufficient to find two messages $M, M' \in \mathcal{M}$, and a ciphertext $C \in \mathcal{C}$, such that $\Pr[\mathsf{Enc}_K(M) = C] \neq \Pr[\mathsf{Enc}_K(M') = C]$.

To this end, let $C \in \mathcal{C}$ be an arbitrary ciphertext, and let $\mathcal{M}(C)$ denote the set of all messages which are possible decryptions of C; that is:

$$\mathcal{M}(C) = \{M \mid M = \mathsf{Dec}_K(C) \text{ for some } K \in \mathcal{K}\}.$$

In other words: $M \in \mathcal{M}(C)$ if there exists some key $K \in \mathcal{K}$, such that $M = \mathsf{Dec}_K(C)$.

First, let M be an arbitrary message in the set $\mathcal{M}(C)$. Hopefully it should be clear that $|\mathcal{M}(C)| \leq |\mathcal{K}|$, because for each distinct message $M \in \mathcal{M}(C)$ there corresponds one or more keys in $\mathcal{K}$. However, since we assumed that $|\mathcal{K}| < |\mathcal{M}|$ we also have

$$|\mathcal{M}(C)| < |\mathcal{M}|.$$

This means there must *exist* some $M' \in \mathcal{M}$ which is *not* in $\mathcal{M}(C)$. In particular, $M' \neq M \in \mathcal{M}(C)$.

Now let us calculate the probabilities $\Pr[\mathsf{Enc}_K(M) = C]$ and $\Pr[\mathsf{Enc}_K(M') = C]$. For the first one we have $\Pr[\mathsf{Enc}_K(M) = C] = p$ for some probability $p > 0$ (we don't actually care what p is, as long as it's non-zero). For the second one, since we explicitly chose M' *not* to be in the set $\mathcal{M}(C)$, we of course have $\Pr[\mathsf{Enc}_K(M') = C] = 0$. Thus

$$\Pr[\mathsf{Enc}_K(M) = C] \neq \Pr[\mathsf{Enc}_K(M') = C],$$

which proves the theorem. $\square$