Appendix

Proof of Lemma 1

Lemma 1. $u \in \Gamma(\pi_{\mathbf{x}})$ *if and only if* $x_u a_u k_u L > c_u$.

Proof. First we show that $V^*(s) \geq 0$ ($\forall s \in \mathcal{S}$):

$$V^*(s) = \arg\max_{a \in \mathcal{A}^s} Q(s,a) \geq Q(s, a = stop) = 0.$$

If direction: Consider state $s = \{u\}$, we have

$$\begin{aligned} V^*(s) &= (1-x_u)(V^*(s) - c_u) + a_u x_u k_u (L - c_u) \\ &\quad + x_u(1 - a_u k_u)(V^*(s^{-u}) - c_u) \\ &\geq (1-x_u)(V^*(s) - c_u) + a_u x_u k_u (L - c_u) \\ &\quad + x_u(1 - a_u k_u)(0 - c_u). \end{aligned}$$

If $x_u a_u k_u L > c_u$, then $V^*(s) > 0$, which means that s is a reachable state and the optimal action at state s is to attack user u instead of stop attacking. Therefore, u belongs to the potential attack set $\Gamma(\pi_{\mathbf{x}})$.

Only if direction: First, consider state s and s^{-u}. If we restrict the attacker's policy so that he never attacks u, then s and s^{-u} are indifferent so that $V^*(s){=}V^*(s^{-u})$. Without the restriction, we have $V^*(s){\geq}V^*(s^{-u})$. In other words, adding a user to a state does not decrease its value. We prove that if $\pi_{\mathbf{x}}(s){=}u$, then $x_u > \frac{c_u}{L a_u k_u}$. By definition we have:

$$\begin{aligned} V^*(s) &= (1-x_u)(V^*(s) - c_u) + a_u x_u k_u (L - c_u) \\ &\quad + x_u(1 - a_u k_u)(V^*(s^{-u}) - c_u). \end{aligned}$$

By adjusting the terms we have:

$$V^*(s) = -\frac{c_u}{a_u x_u} + L k_u + (1 - k_u) V^*(s^{-u}).$$

Since $V^*(s){\geq}V^*(s^{-u})$, then:

$$-\frac{c_u}{a_u x_u} + L k_u \geq k_u V^*(s^{-u}) \geq 0$$

Note that if $-\frac{c_u}{a_u x_u}{+}Lk_u{=}0$, we have $V^*(s){=}V^*(s^{-u}){=}0$ and $s{=}\{u\}$. Due to the setting that the attacker always prefers stopping attack rather than launching another attack, we have $\pi_{\mathbf{x}}(s){=}0$, which contradicts the assumption that $\pi_{\mathbf{x}}(s){=}u$. Therefore, $-\frac{c_u}{a_u x_u} + Lk_u > 0$, equivalently, $x_u > \frac{c_u}{L a_u k_u}$. □

Proof of Lemma 2

Lemma 2.

$$\theta(\mathbf{x}, \pi_{\mathbf{x}}) = \begin{cases} 1 - \prod_{u \in \Gamma(\pi_{\mathbf{x}})}(1 - a_u k_u), & \textit{if } \Gamma(\pi_{\mathbf{x}}) \neq \emptyset \\ 0, & \textit{if } \Gamma(\pi_{\mathbf{x}}) = \emptyset \end{cases}.$$

Proof. If $\Gamma(\pi_{\mathbf{x}}) = \emptyset$, meaning that the attacker stops attacking at the initial state s_0, therefore the probability that the credential accessed is 0. Otherwise, we write the reachable states set as $\Delta(\pi_{\mathbf{x}}) = \{s_0, s_1, ..., s_r\} \cup \{s^n, s^y\}$. We denote by $M_{\Delta(\pi_{\mathbf{x}})}$ the transition probability matrix, whose entry M_{ij} represents the probability that state s_i transitions to s_j under policy $\pi_{\mathbf{x}}$ (WLOG, we define $s_{r+1}{=}s^n$ and $s_{r+2}{=}s^y$). There are two cases for s_r: (1) $\pi_{\mathbf{x}}(s_r) = u \in \mathcal{A}^{s_r}$ and (2) $\pi_{\mathbf{x}}(s_r) = stop$.

If case (1), s_r could transition to itself, s^n or s^y. Hence $M_{\Delta(\pi_{\mathbf{x}})}$ has the form like (denote $d_i{=}a_{u^i}k_{u^i}$ and $x_i = x_{u^i}$):

$$\begin{bmatrix} 1{-}x_0 & x_0(1{-}d_0) & & & & & d_0 x_0 \\ & 1{-}x_1 & x_1(1{-}d_1) & & & & d_1 x_1 \\ & & \ddots & \ddots & & & \vdots \\ & & & 1{-}x_r & x_r(1{-}d_r) & & d_r x_r \\ & & & & 1 & & \\ & & & & & & 1 \end{bmatrix}$$

Precisely, $M_{\Delta(\pi_{\mathbf{x}})}$ can be represented as:

$$M_{\Delta(\pi_{\mathbf{x}})} = \left[\begin{array}{c:c} A & B \\ \hdashline 0 & I_2 \end{array}\right]$$

where A is $r{+}1$ dimensional square matrix, I_2 is 2 dimensional unit diagonal matrix and B is $(r{+}1) \times 2$ matrix. We introduce a $(r{+}1) \times 2$ matrix E:

$$E = FB, \text{ where } F = (I_{r+1} - A)^{-1}$$

Note that s^n and s^y are absorbing states. According to the properties of absorbing Markov chain, s_0 will eventually end in state s^n or s^y with probability E_{11} and E_{12} respectively, and $E_{11}{+}E_{12}{=}1$. Therefore, the probability of losing the credential is equal to the probability that the attacker eventually ends in state s^y, i.e., $\theta(\mathbf{x}, \pi_{\mathbf{x}}){=}E_{12}$. We can directly calculate E_{11} based on the rules of matrix calculation:

$$\begin{aligned} E_{11} &= \sum_{i=1}^{r+1} F_{1i} B_{i1} \\ &= F_{1,r+1} B_{r+1,1} \\ &= \frac{\prod_{i=0}^{r-1}(1 - d_i)}{x_r} x_r (1 - d_r) \\ &= \prod_{i=0}^{r} (1 - d_i) \\ &= \prod_{u \in \Gamma(\pi_{\mathbf{x}})} (1 - a_u k_u) \end{aligned}$$

Then $E_{12} = 1 - E_{11} = 1 - \prod_{u \in \Gamma(\pi_{\mathbf{x}})}(1 - a_u k_u)$.

If case (2), s_r transitions to s^n with probability 1. Thus $M_{\Delta(\pi_{\mathbf{x}})}$ has the form like ($d_i{=}a_{u^i}k_{u^i}$ and $x_i = x_{u^i}$):

$$
\begin{bmatrix}
1-x_0 & x_0(1-d_0) & & & & & x_0k_0 \\
& 1-x_1 & x_1(1-d_1) & & & & x_1k_1 \\
& & \ddots & \ddots & & & \vdots \\
& & & 1-x_{r-1} & x_{r-1}(1-d_{r-1}) & 0 & x_{r-1}d_{r-1} \\
& & & & & 1 & \\
& & & & & 1 & \\
& & & & & & 1
\end{bmatrix}
$$

Similarly,

$$
\begin{aligned}
E_{11} &= \sum_{i=1}^{r+1} F_{1i}B_{i1} \\
&= F_{1,r+1} \\
&= \prod_{i=0}^{r-1}(1-d_i) \\
&= \prod_{u\in\Gamma(\pi_{\mathbf{x}})} (1-a_uk_u)
\end{aligned}
$$

Then, we still have $E_{12} = 1 - E_{11} = 1 - \prod_{u\in\Gamma(\pi_{\mathbf{x}})}(1 - a_uk_u)$. □

Proof of Theorem 1

Theorem 1. *The defender's expected utility remains the same no matter how the attacker breaks ties, i.e., choosing any optimal policy.*

Proof. Recall that in single-credential case the defender's utility function is

$$P_d(\mathbf{x}, \pi_{\mathbf{x}}) = -\rho^{\mathcal{T}}\theta(\mathbf{x}, \pi_{\mathbf{x}})L - \sum_{u\in U} \Lambda(x_u).$$

Based on the result of Lemma 1, $\Gamma(\pi_{\mathbf{x}})$ can be represented as $\{u \in U | x_u > \frac{c_u}{La_uk_u}\}$, then $\theta(\mathbf{x}, \pi_{\mathbf{x}})$ can be represented as

$$\theta(\mathbf{x}, \pi_{\mathbf{x}}) = 1 - \prod_{u\in\{u'\in U | x_{u'} > \frac{c_{u'}}{La_{u'}k_{u'}}\}} (1-k_u).$$

For any other optimal policy $\pi'_{\mathbf{x}}$, we have

$$\theta(\mathbf{x}, \pi'_{\mathbf{x}}) = 1 - \prod_{u\in\{u'\in U | x_{u'} > \frac{c_{u'}}{La_{u'}k_{u'}}\}} (1-k_u).$$

Note that $\theta(\mathbf{x}, \pi_{\mathbf{x}}) = \theta(\mathbf{x}, \pi_{\mathbf{x}})'$, which indicates that the defender's expected utility will be the same when the attacker chooses any other optimal policy. □

Proof of Theorem 2

Theorem 2. *x_u^1 is an arbitrary point in $\arg\min_{x\in[0,\frac{c_u}{La_uk_u}]} \Lambda_u$ and x_u^2 is an arbitrary point in $\arg\min_{x\in(\frac{c_u}{La_uk_u},1]} \Lambda_u$.*

Proof. Recall that in single-credential case the defender's utility function is

$$P_d(\mathbf{x}, \pi_{\mathbf{x}}) = -\rho^{\mathcal{T}}\theta(\mathbf{x}, \pi_{\mathbf{x}})L - \sum_{u\in U} \Lambda_u(x_u).$$

Consider a user u, given all values of $x_{u'}$ ($u' \in U \setminus \{u\}$), $\theta(\mathbf{x}, \pi_{\mathbf{x}})$ is constant for any $x_u \in [0, \frac{c_u}{La_uk_u}]$ since the potential attack set $\Gamma(\pi_{\mathbf{x}})$ remains the same when x_u varies among $[0, \frac{c_u}{La_uk_u}]$. Therefore, any point in $\arg\min_{x\in[0,\frac{c_u}{La_uk_u}]} \Lambda_u$ maximizes $P_d(\mathbf{x}, \pi_{\mathbf{x}})$. Similarly, $\theta(\mathbf{x}, \pi_{\mathbf{x}})$ is constant for any $x_u \in (\frac{c_u}{La_uk_u}, 1]$. Therefore, any points in $\arg\min_{x\in(\frac{c_u}{La_uk_u},1]} \Lambda_u$ maximizes $P_d(\mathbf{x}, \pi_{\mathbf{x}})$. □