(a) X_f is not a resource monotone A (state) monotone for this resource theory would have to satisfy, for all free local channels $\mathcal{T}_A, \mathcal{T}_B$ and for the fixed allowed POVMs M_A, M_B,

$$X_f(\rho_{ABC}; M_A, M_B) \;\geq\; X_f\big((\mathcal{T}_A \otimes \mathrm{id}_{BC})(\rho_{ABC}); M_A, M_B\big),$$

and similarly for $\mathcal{T}_B$. We show that unless f is constant, there exist allowed choices of $(\rho_{ABC}, M_A, M_B, \mathcal{T}_A)$ for which X_f *increases* under a free operation, hence X_f is not a monotone.

To see this, take A, B, C to be two-dimensional with computational bases $\{|0\rangle, |1\rangle\}$. Let ρ_{ABC} be the zero state, and $\mathcal{T}_A$ a classical permutation channel (bit-flip) on A

$$\rho_{ABC} = |0,0,0\rangle\langle 0,0,0|, \qquad \mathcal{T}_A(|0\rangle\langle 0|) = |1\rangle\langle 1|, \quad \mathcal{T}_A(|1\rangle\langle 1|) = |0\rangle\langle 0|.$$

Then $\rho'_{ABC} = (\mathcal{T}_A \otimes \mathrm{id}_{BC})(\rho_{ABC}) = |1,0,0\rangle\langle 1,0,0|$.

Choose any allowed diagonal POVM on B such that $\mathrm{Tr}(M_B^0 |0\rangle\langle 0|) > 0$; for instance take $o(0|0) = o(0|1) = \frac{1}{2}$. On A, choose an allowed diagonal POVM with (say) $m(0|0) = \alpha$ and $m(0|1) = \beta$, where $\alpha, \beta \in (0,1)$. (Other entries are fixed by normalization: $m(1|a) = 1 - m(0|a)$.)

Because ρ_{ABC} and ρ'_{ABC} are basis states, the sum defining X_f collapses to a single term:

$$X_f(\rho_{ABC}; M_A, M_B) = f\left(\frac{\mathrm{Tr}(M_A^0 |0\rangle\langle 0|)}{\mathrm{Tr}(M_B^0 |0\rangle\langle 0|)}\right) = f\left(\frac{\alpha}{o(0|0)}\right), \qquad X_f(\rho'_{ABC}; M_A, M_B) = f\left(\frac{\mathrm{Tr}(M_A^0 |1\rangle\langle 1|)}{\mathrm{Tr}(M_B^0 |0\rangle\langle 0|)}\right) = f\left(\frac{\beta}{o(0|0)}\right).$$

If f is not constant, we can find α, β so that $f(\beta/o(0|0)) > f(\alpha/o(0|0))$. This yields $X_f(\rho'_{ABC}; M_A, M_B) > X_f(\rho_{ABC}; M_A, M_B)$, i.e. X_f increases under the free operation $\mathcal{T}_A$. Therefore X_f is not a valid resource monotone (except in the trivial case where f is a constant function).

(b) Operational interpretation. *Resource theory.* The allowed maps $\mathcal{T}_A, \mathcal{T}_B$ are exactly *classical stochastic post-processings* of a and b, and the allowed POVMs $\{M_A^c\}_c, \{M_B^c\}_c$ are diagonal, i.e. they implement conditional distributions $m(c|a)$ and $o(c|b)$ producing a classical output c. Thus the resource theory describes what can be done using only *local classical processing and classical measurements* in the fixed bases.

Proposed measure. For a state with diagonal $p(a,b,c) = \langle a,b,c|\rho_{ABC}|a,b,c\rangle$,

$$X_f(\rho_{ABC}; M_A, M_B) = \mathbb{E}_{(a,b,c)\sim p}\left[f\left(\frac{m(c|a)}{o(c|b)}\right)\right]$$

is the expected value of a score comparing Alice's and Bob's likelihoods assigned to the realised label c (i.e. a relative "Alice-versus-Bob" prediction payoff determined by f). To gain intuition, consider $f(x) = \log x$ and write $p(a,b,c) = \langle a,b,c|\rho_{ABC}|a,b,c\rangle$ for a state diagonal in this basis. For generic allowed POVMs (i.e. generic conditionals) $m(c|a)$ and $o(c|b)$,

$$\begin{aligned} X_{\log}(\rho_{ABC}; M_A, M_B) &= \sum_{a,b,c} p(a,b,c)\, \log \frac{m(c|a)}{o(c|b)} \\ &= \sum_{a,c} p(a,c) \log m(c|a) \;-\; \sum_{b,c} p(b,c) \log o(c|b). \end{aligned}$$

Add and subtract the true conditionals to obtain the decomposition

$$\sum_{a,c} p(a,c) \log m(c|a) = -H(C|A) - \sum_a p(a)\, D\big(p(\cdot|a) \,\|\, m(\cdot|a)\big),$$

Therefore

$$X_{\log}(\rho_{ABC}; M_A, M_B) = \big(H(C|B) - H(C|A)\big) + \sum_b p(b)\, D\big(p(\cdot|b) \,\|\, o(\cdot|b)\big) - \sum_a p(a)\, D\big(p(\cdot|a) \,\|\, m(\cdot|a)\big),$$

i.e.

$$X_{\log} = I(A;C) - I(B;C) \;+\; (\text{Bob penalty}) \;-\; (\text{Alice penalty}).$$

In particular, if we choose $m(c|a) = p(c|a)$ and $o(c|b) = p(c|b)$, the relative entropy terms vanish and $X_{\log}(\rho_{ABC}; M_A, M_B) = H(C|B) - H(C|A) = I(A;C) - I(B;C)$. This makes the interpretation in terms of an Alice versus Bob advantage transparent: $H(C|A)$ is the remaining uncertainty about the label C for a player who knows A, and $H(C|B)$ is the remaining uncertainty for a player who knows B. Thus X_f quantifies *Alice's advantage over Bob at predicting/guessing C.*

Marks were awarded for clear, rigorous arguments (correctly stating the monotonicity requirement and either proving it or giving a valid counterexample), and for demonstrating conceptual understanding by justifying any interpretation with precise mathematical statements. We rewarded solutions that were both mathematically correct and explanatory: full credit required a logically complete argument plus a brief justification of the underlying intuition in formal terms (rather than unsupported claims such as "it's Alice vs Bob").

Side remark. Here's a neat argument to see that X_f is not a good monotone for $f(x) = log(x)$. Let $A \to A'$ be any free local channel (a stochastic map in the fixed basis). By the data processing inequality,

$$I(A'; C) \leq I(A; C),$$

hence $I(A; C) - I(B; C)$ cannot increase under free processing on A. However, for a free local channel on B we have similarly $I(B'; C) \leq I(B; C)$, and therefore

$$I(A; C) - I(B'; C) \ \geq \ I(A; C) - I(B; C),$$

i.e. the same expression *increases* whenever Bob applies a channel to his information about C.

Now let $A \to A'$ be any free local channel (a stochastic map in the fixed basis). By the data processing inequality,

$$I(A'; C) \leq I(A; C),$$

hence $I(A; C) - I(B; C)$ cannot increase under free processing on A. However, for a free local channel on B we have similarly $I(B'; C) \leq I(B; C)$, and therefore

$$I(A; C) - I(B'; C) \ \geq \ I(A; C) - I(B; C),$$

i.e. the same expression *increases* whenever Bob applies a channel to his information about C. Since the resource theory declares *both* A-side and B-side stochastic maps free, this already shows that $X_{\log(x)}$ cannot simultaneously be a valid monotone under the two types of stated free operations. Therefore $X_{\log(x)}$ is not a resource monotone.

(c) Bonus: what about $g(X_f)$? If X_f fails to be monotone, then the composition with another function g also fails in general. Only very special choices of functions can avoid this conclusion.

We have already seen that there exist free operations ε and states ρ_{ABC} such that

$$X_f(\varepsilon(\rho_{ABC}); M_A, M_B) \ > \ X_f(\rho_{ABC}; M_A, M_B).$$

But that there also exist operations and states for which the opposite happens (recall the log example, where one can also find free operations and states for which the *opposite* strict inequality holds, depending on whether Alice's or Bob's system is processed.) Applying g to both sides:

- If g is **increasing**, it preserves the order, hence

$$g\big(X_f(\varepsilon(\rho_{ABC}); M_A, M_B)\big) \ > \ g\big(X_f(\rho_{ABC}; M_A, M_B)\big),$$

  so $g\big(X_f(\cdot; M_A, M_B)\big)$ is also *not* a monotone.

- If g is **decreasing**, it reverses the order, hence

$$g\big(X_f(\varepsilon(\rho_{ABC}); M_A, M_B)\big) \ < \ g\big(X_f(\rho_{ABC}; M_A, M_B)\big).$$

  Thus composing with a function g does not "repair" monotonicity.

- If g is the **identity** map, then $g(X_f) = X_f$, so the same violation remains.

- If g is constant on the relevant range, it yields a trivial quantity.

To construct a monotone which makes sense for the structure specified by the problem, we need stringent conditions. Consider first the following example. Let d_C be any function on quantum states that is *contractive* under CPTP maps, i.e.

$$d_C\big(\varepsilon(\rho), \varepsilon(\sigma)\big) \leq d_C(\rho, \sigma) \qquad \forall\, \rho, \sigma, \ \forall\, \text{CPTP } \varepsilon.$$

Given a resource theory with free set $\mathcal{F}$, define the distance-to-free-set monotone

$$M_C(\rho) := \inf_{\gamma \in \mathcal{F}} d_C(\rho, \gamma).$$

If ε is a free operation, then $\varepsilon(\mathcal{F}) \subseteq \mathcal{F}$, hence

$$\begin{aligned} M_C(\varepsilon(\rho)) &= \inf_{\gamma\in\mathcal{F}} d_C\big(\varepsilon(\rho),\gamma\big) \\ &\le \inf_{\gamma\in\mathcal{F}} d_C\big(\varepsilon(\rho),\varepsilon(\gamma)\big) \\ &\le \inf_{\gamma\in\mathcal{F}} d_C(\rho,\gamma) = M_C(\rho), \end{aligned}$$

where the first inequality uses $\varepsilon(\mathcal{F}) \subseteq \mathcal{F}$ and the second uses contractivity. Thus M_C is monotone under all free operations. An example of this is the quantum relative entropy, in which case contractivity becomes the data processing inequality.

In our setting the free CPTP maps are *local* and act only on one register at a time, i.e.

$$\varepsilon_A = \mathcal{T}_A \otimes \mathrm{id}_{BC}, \qquad \varepsilon_B = \mathrm{id}_{AC} \otimes \mathcal{T}_B,$$

with $\mathcal{T}_A, \mathcal{T}_B$ stochastic (classical) channels. Hence, to use the "distance-to-free-set" construction it is enough to pick a function d that is *separately contractive* under these local maps:

$$d(\varepsilon_A(\rho),\varepsilon_A(\sigma)) \le d(\rho,\sigma), \qquad d(\varepsilon_B(\rho),\varepsilon_B(\sigma)) \le d(\rho,\sigma) \quad \forall\, \rho,\sigma.$$

Then, for $M(\rho) := \inf_{\gamma\in\mathcal{F}} d(\rho,\gamma)$ and any free ε_A (and similarly ε_B),

$$M(\varepsilon_A(\rho)) = \inf_{\gamma\in\mathcal{F}} d(\varepsilon_A(\rho),\gamma) \le \inf_{\gamma\in\mathcal{F}} d(\varepsilon_A(\rho),\varepsilon_A(\gamma)) \le \inf_{\gamma\in\mathcal{F}} d(\rho,\gamma) = M(\rho),$$

using $\varepsilon_A(\mathcal{F}) \subseteq \mathcal{F}$ and the corresponding contractivity. Thus M is monotone under free processing on A and on B.