

## A Proofs of Propositions

To simplify our notation, we remove subscripts such that  $X, A, Y$  denotes the transition  $X_t, A_t, X_{t+1}$ , and  $Z$  denotes the latent  $Z_{t+1}$ . Then the environment's dynamics is given by  $\tau(Y|x, a)$ , the agent's policy is given by  $\pi(A|x)$ , and the induced state visitation given by  $\rho_\pi(X)$ . The generator is denoted  $p_\theta(Z|x, a, y)$ , the reconstructor  $f_\eta(x, a, z)$ , and the critic  $g_\nu(x, a, z)$ . We start with several lemmas that will be useful, the first being a pointwise version of Barber and Agakov's variational lower bound:

**Lemma 4 (Pointwise Barber-Agakov)** Denote the pointwise mutual information:

$$\text{PMI}_\theta(x, a; z) := \log \frac{p_\theta(z|x, a)}{p_\theta(z)} \quad (20)$$

Then for any variational distribution  $q$ :

$$\mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \text{PMI}_\theta(x, a; Z) \geq \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \log \frac{q(Z|x, a)}{p_\theta(Z)} \quad (21)$$

*Proof.* Starting from the left hand side:

$$\mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \text{PMI}_\theta(x, a; Z) = \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \log \frac{p_\theta(Z|x, a)}{p_\theta(Z)} \quad (22)$$

$$= \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \log \frac{p_\theta(Z|x, a)}{p_\theta(Z)} + \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \log \frac{q(Z|x, a)}{q(Z|x, a)} \quad (23)$$

$$= \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \log \frac{q(Z|x, a)}{p_\theta(Z)} + D_{\text{KL}}(p_\theta(Z|x, a) \| q(Z|x, a)) \quad (24)$$

$$\geq \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \log \frac{q(Z|x, a)}{p_\theta(Z)} \quad (25)$$

which completes the proof.  $\square$

Next, we define a generic contrastive expression with  $K - 1$  “negative” samples of  $Z$ , and show that taking its expectation with respect to those samples yields a valid (i.e. normalized) probability density:

**Lemma 5 (Normalized Variational)** Given independent samples  $z_{1:K-1}$  from  $p_\theta$ , define:

$$q(z|x, a, z_{1:K-1}) := \frac{p_\theta(z) \cdot e^{g_\nu(x, a, z)}}{\frac{1}{K} \left( e^{g_\nu(x, a, z)} + \sum_{i=1}^{K-1} e^{g_\nu(x, a, z_i)} \right)} \quad (26)$$

then the following defines a normalized density:

$$q(Z|x, a) := \mathbb{E}_{Z_{1:K-1} \sim p_\theta^{K-1}} q(Z|x, a, Z_{1:K-1}) \quad (27)$$

*Proof.* The expectation integrates to one:

$$\int_Z q(z|x, a) dz = \int_Z \mathbb{E}_{Z_{1:K-1} \sim p_\theta^{K-1}} \frac{p_\theta(z) \cdot e^{g_\nu(x, a, z)}}{\frac{1}{K} \left( e^{g_\nu(x, a, z)} + \sum_{i=1}^{K-1} e^{g_\nu(x, a, z_i)} \right)} dz \quad (28)$$

$$= \mathbb{E}_{\substack{Z \sim p_\theta \\ Z_{1:K-1} \sim p_\theta^{K-1}}} \frac{e^{g_\nu(x, a, Z)}}{\frac{1}{K} \left( e^{g_\nu(x, a, Z)} + \sum_{i=1}^{K-1} e^{g_\nu(x, a, Z_i)} \right)} \quad (29)$$

$$= K \cdot \mathbb{E}_{Z_{1:K} \sim p_\theta^K} \frac{e^{g_\nu(x, a, Z_1)}}{\sum_{i=1}^K e^{g_\nu(x, a, Z_i)}} \quad (30)$$

$$= \mathbb{E}_{Z_{1:K} \sim p_\theta^K} \frac{\sum_{j=1}^K e^{g_\nu(x, a, Z_j)}}{\sum_{i=1}^K e^{g_\nu(x, a, Z_i)}} = 1 \quad (31)$$

which completes the proof.  $\square$

These two results allow us to show that the information  $Z$  contains on a tuple  $x, a$ —with respect to the generator parameterized as  $\theta$ —is lower-bounded by the  $x, a$ -conditioned contrastive loss between “positive” samples  $Z \sim p_\theta(\cdot|x, a)$  from the posterior and “negative” samples  $Z \sim p_\theta$  from the prior:

**Lemma 6 (State-Action Lower Bound)** The  $x, a$ -wise mutual information satisfies:

$$\begin{aligned} & \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \text{PMI}_\theta(x, a; Z) \\ & \geq \mathbb{E}_{\substack{Z \sim p_\theta(\cdot|x, a) \\ Z_{1:K-1} \sim p_\theta^{K-1}}} \log \frac{e^{g_\nu(x, a, Z)}}{\frac{1}{K} \left( e^{g_\nu(x, a, Z)} + \sum_{i=1}^{K-1} e^{g_\nu(x, a, Z_i)} \right)} \end{aligned} \quad (32)$$

*Proof.* Use Lemmas 4 and 5, then Jensen’s inequality:

$$\begin{aligned} & \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \text{PMI}_\theta(x, a; Z) \\ & \geq \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \log \frac{q(Z|x, a)}{p_\theta(Z)} \end{aligned} \quad (33)$$

$$= \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \log \mathbb{E}_{\substack{Z_{1:K-1} \sim p_\theta^{K-1} \\ p_\theta(Z)}} \frac{q(Z|x, a, Z_{1:K-1})}{p_\theta(Z)} \quad (34)$$

$$\geq \mathbb{E}_{\substack{Z \sim p_\theta(\cdot|x, a) \\ Z_{1:K-1} \sim p_\theta^{K-1}}} \log \frac{q(Z|x, a, Z_{1:K-1})}{p_\theta(Z)} \quad (35)$$

$$= \mathbb{E}_{\substack{Z \sim p_\theta(\cdot|x, a) \\ Z_{1:K-1} \sim p_\theta^{K-1}}} \log \frac{e^{g_\nu(x, a, Z)}}{\frac{1}{K} \left( e^{g_\nu(x, a, Z)} + \sum_{i=1}^{K-1} e^{g_\nu(x, a, Z_i)} \right)} \quad (36)$$

which completes the proof.  $\square$

Next, we show that our invariance loss (Objective 2) for a tuple  $x, a, z$  is equal to the pointwise mutual information in the limit of infinitely large negative batches, assuming an optimal critic parameter:

**Lemma 7 (Pointwise Asymptotic Equality)** Define the pointwise invariance loss:

$$\mathcal{L}_{\theta, \nu}^K(x, a, z) := \mathbb{E}_{Z_{1:K-1} \sim p_\theta^{K-1}} \log \frac{e^{g_\nu(x, a, z)}}{\frac{1}{K} \left( e^{g_\nu(x, a, z)} + \sum_{i=1}^{K-1} e^{g_\nu(x, a, Z_i)} \right)} \quad (37)$$

and the optimal critic parameter:

$$\nu^* := \arg \max_{\nu} \mathbb{E}_{\substack{X \sim \rho_\pi \\ A \sim \pi(\cdot|X) \\ Y \sim \tau(\cdot|X, A) \\ Z \sim p_\theta(\cdot|X, A, Y)}} \mathcal{L}_{\theta, \nu}^K(X, A, Z) \quad (38)$$

Then  $\lim_{K \rightarrow \infty} \mathcal{L}_{\theta, \nu^*}^K(x, a, z) = \text{PMI}_\theta(x, a; z)$ .

*Proof.* The  $\mathbb{E}[\mathcal{L}_{\theta, \nu}^K(X, A, Z)]$  term is just the InfoNCE loss between variables  $Z$  and  $X, A$ , so we know that  $\nu^*$  satisfies  $g_{\nu^*}(x, a, z) = \log \frac{p_\theta(z|x, a)}{p_\theta(z)} + c(x, a)$ . Substituting this back into  $\mathcal{L}_{\theta, \nu}^K(x, a, z)$ :

$$\lim_{K \rightarrow \infty} \mathcal{L}_{\theta, \nu^*}^K(x, a, z) \quad (39)$$

$$= \lim_{K \rightarrow \infty} \mathbb{E}_{Z_{1:K-1} \sim p_\theta^{K-1}} \log \frac{e^{g_{\nu^*}(x, a, z)}}{\frac{1}{K} \left( e^{g_{\nu^*}(x, a, z)} + \sum_{i=1}^{K-1} e^{g_{\nu^*}(x, a, Z_i)} \right)} \quad (40)$$

$$= \lim_{K \rightarrow \infty} \mathbb{E}_{Z_{1:K-1} \sim p_\theta^{K-1}} \log \frac{\frac{p_\theta(z|x, a)}{p_\theta(z)}}{\frac{1}{K} \left( \frac{p_\theta(z|x, a)}{p_\theta(z)} + \sum_{i=1}^{K-1} \frac{p_\theta(Z_i|x, a)}{p_\theta(Z_i)} \right)} \quad (41)$$

$$= \lim_{K \rightarrow \infty} \mathbb{E}_{Z_{1:K-1} \sim p_\theta^{K-1}} \left[ \log \frac{p_\theta(z|x, a)}{p_\theta(z)} - \log \frac{\frac{p_\theta(z|x, a)}{p_\theta(z)} + \sum_{i=1}^{K-1} \frac{p_\theta(Z_i|x, a)}{p_\theta(Z_i)}}{K} \right] \quad (42)$$

$$= \log \frac{p_\theta(z|x, a)}{p_\theta(z)} - \lim_{K \rightarrow \infty} \log \frac{\frac{p_\theta(z|x, a)}{p_\theta(z)} + K - 1}{K} = \text{PMI}_\theta(x, a; z) \quad (43)$$

which completes the proof.  $\square$

This gives us what we need to derive Proposition 1, which we restate using our subscript-less notation:

**Proposition 8 (Optimal Invariance)** The (state-action) invariance bonus satisfies:

$$\mathcal{R}_{\theta, \nu}^{K, \text{inv.}}(x, a) \leq \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \text{PMI}_\theta(x, a; Z) \quad (44)$$

and for the optimal  $\nu^*$  the bound is asymptotically tight as  $K \rightarrow \infty$ :

$$\lim_{K \rightarrow \infty} \mathcal{R}_{\theta, \nu^*}^{K, \text{inv.}}(x, a) = \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \text{PMI}_\theta(x, a; Z) \quad (45)$$

*Proof.* Use Lemma 6 for the first part:

$$\mathcal{R}_{\theta, \nu}^{K, \text{inv.}}(x, a) := \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \mathcal{L}_{\theta, \nu}^K(x, a, Z) \quad (46)$$

$$= \mathbb{E}_{\substack{Z \sim p_\theta(\cdot|x, a) \\ Z_{1:K-1} \sim p_\theta^{K-1}}} \log \frac{e^{g_\nu(x, a, Z)}}{\frac{1}{K} \left( e^{g_\nu(x, a, Z)} + \sum_{i=1}^{K-1} e^{g_\nu(x, a, Z_i)} \right)} \quad (47)$$

$$\leq \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \text{PMI}_\theta(x, a; Z) \quad (48)$$

and use Lemma 7 for the second part:

$$\lim_{K \rightarrow \infty} \mathcal{R}_{\theta, \nu^*}^{K, \text{inv.}}(x, a) = \lim_{K \rightarrow \infty} \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \mathcal{L}_{\theta, \nu^*}^K(x, a, Z) \quad (49)$$

$$= \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \lim_{K \rightarrow \infty} \mathcal{L}_{\theta, \nu^*}^K(x, a, Z) \quad (50)$$

$$= \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \text{PMI}_\theta(x, a; Z) \quad (51)$$

which completes the proof.  $\square$

Next, we show that Proposition 2 is true, which we similarly restate using our subscript-less notation:

**Proposition 9 (Optimal Reconstruction)** Denote with  $\mathcal{R}_{\theta, \eta}^{\text{rec.}}(x, a)$  the reconstruction bonus as in Objective 1,  $\mathcal{R}_{\theta, \nu}^{K, \text{inv.}}(x, a)$  the invariance bonus as in Objective 2, and their weighted sum for any  $\lambda$ :

$$J(\theta, \eta, \nu; \lambda) := \mathbb{E}_{\substack{X \sim \rho_\pi \\ A \sim \pi(\cdot|X)}} \left[ \frac{1}{\lambda} \mathcal{R}_{\theta, \eta}^{\text{rec.}}(X, A) + \lim_{K \rightarrow \infty} \mathcal{R}_{\theta, \nu}^{K, \text{inv.}}(X, A) \right] \quad (52)$$

Then its minimax optimal value is zero:

$$\min_{\theta, \eta} \max_{\nu} J(\theta, \eta, \nu; \lambda) = 0 \quad (53)$$

*Proof.* Take any MDP, as in Figure 1(a). By reparameterization, we know there exists an equivalent graphical representation under which  $Z$  is exogenous, as in Figure 1(b). Assuming realizability, let  $\eta^*$  be such that  $f_{\eta^*} = f$ , and let  $\theta^*$  be such that  $p_{\theta^*}(Z|x, a, y) = p_{\eta^*}(Z|x, a, y)$  for any  $x, a, y$ . First, by construction we have that  $Z \perp X, A$ , so the mutual information between  $Z$  and  $X, A$  must be zero:

$$\mathbb{E}_{\substack{X \sim \rho_\pi \\ A \sim \pi(\cdot|X)}} \left[ \lim_{K \rightarrow \infty} \mathcal{R}_{\theta, \nu^*}^{K, \text{inv.}}(X, A) \right] = \mathbb{E}_{\substack{X \sim \rho_\pi \\ A \sim \pi(\cdot|X) \\ Z \sim p_\theta(\cdot|X, A)}} \text{PMI}_\theta(X, A; Z) \quad (54)$$

$$= \mathbb{I}_\theta[X, A; Z] = 0 \quad (55)$$

for optimal critic parameter  $\nu^*$ , where the first equality uses Proposition 1. Second, by consistency of counterfactuals  $f_{\eta^*}(x, a, Z) = y$  for any  $Z \sim p_{\theta^*}(\cdot|x, a, y)$ , so the reconstruction term is also zero. It is easy to verify the optimal critic is a maximizer, and the optimal generator/reconstructor minimizers, which completes the proof.  $\square$

Finally, we recall the following basic relationship:

**Lemma 10 (Conditional Mutual Information)** Conditioned on any  $x, a$ , we have that:

$$\mathbb{I}_\theta[Y; Z|x, a] = \mathbb{H}[Y|x, a] + \mathbb{H}_\theta[Y|x, a, Z] \quad (56)$$

*Proof.* Starting from the left hand side:

$$\mathbb{I}_\theta[Y; Z|x, a] := \mathbb{E}_{Z \sim p_\theta} D_{\text{KL}}(p_\theta(Y|x, a, Z) \| \tau(Y|x, a)) \quad (57)$$

$$= \mathbb{E}_{\substack{Z \sim p_\theta \\ Y \sim p_\theta(\cdot|x, a, Z)}} \log p_\theta(Y|x, a, Z) - \mathbb{E}_{\substack{Z \sim p_\theta \\ Y \sim p_\theta(\cdot|x, a, Z)}} \tau(Y|x, a) \quad (58)$$

$$= - \int_Z p_\theta(z) \mathbb{H}_\theta[Y|x, a, z] dz - \mathbb{E}_{\substack{Y \sim \tau(\cdot|x, a) \\ Z \sim p_\theta(\cdot|x, a, Y)}} \tau(Y|x, a) \quad (59)$$

$$= \mathbb{H}[Y|x, a] - \mathbb{H}_\theta[Y|x, a, Z] \quad (60)$$

which completes the proof.  $\square$

Now, in our structural causal model, by construction  $Z$  captures all sources of noise—that is, there is no residual noise in each outcome  $Y$ . However, for the purposes of optimization, while learning  $\eta$  we let the residual error be captured by a Gaussian “log-likelihood” (note that  $\lambda$  plays the role of “ $2\sigma^2$ ”):

$$\log p_\eta(Y|x, a, z) := -\frac{1}{2} \log(\lambda\pi) - \frac{1}{\lambda} (Y - f_\eta(x, a, z))^2 \quad (61)$$

and note that  $\theta$  also induces a log-likelihood of the “ground-truth” conditional:

$$\log p_\theta(Y|x, a, z) := \log \frac{p_\theta(z|x, a, Y) \tau(Y|x, a) \pi(a, x) \rho_\pi(x)}{\int_Y p_\theta(z|x, a, y) \tau(y|x, a) \pi(a|x) \rho_\pi(x) dy} \quad (62)$$

Now, recall the reconstruction loss and (state-action) reconstruction bonus:

$$\mathcal{L}_\eta(x, a, z, y) := \left\| y - f_\eta(x, a, z) \right\|_2^2 \quad (63)$$

$$\mathcal{R}_{\theta, \eta}^{\text{rec.}}(x, a) := \mathbb{E}_{\substack{Y \sim \tau(\cdot|x, a) \\ Z \sim p_\theta(\cdot|x, a, Y)}} \mathcal{L}_\eta(x, a, Z, Y) \quad (64)$$

as well as the invariance loss and (state-action) invariance bonus:

$$\mathcal{L}_{\theta, \nu}^K(x, a, z) := \mathbb{E}_{\substack{(X_1, \dots, X_{K-1}) \sim \prod_{i=1}^{K-1} \rho_\pi \\ (A_1, \dots, A_{K-1}) \sim \prod_{i=1}^{K-1} \pi(\cdot|X_i) \\ (Y_1, \dots, Y_{K-1}) \sim \prod_{i=1}^{K-1} \tau(\cdot|X_i, A_i) \\ (Z_1, \dots, Z_{K-1}) \sim \prod_{i=1}^{K-1} p_\theta(\cdot|X_i, A_i, Y_i)}} \log \frac{e^{g_\nu(x, a, z)}}{\frac{1}{K} \left( e^{g_\nu(x, a, z)} + \sum_{i=1}^{K-1} e^{g_\nu(x, a, Z_i)} \right)} \quad (65)$$

$$\mathcal{R}_{\theta, \nu}^{\text{inv.}}(x, a) := \mathbb{E}_{Z \sim p_\theta(\cdot|x, a)} \mathcal{L}_{\theta, \nu}^K(x, a, Z) \quad (66)$$

Moreover, recall the hindsight intrinsic reward function:

$$\mathcal{R}_{\theta, \eta, \nu^*}(x, a) := \frac{1}{\lambda} \mathcal{R}_{\theta, \eta}^{\text{rec.}}(x, a) + \lim_{K \rightarrow \infty} \mathcal{R}_{\theta, \nu^*}^{\text{inv.}}(x, a) \quad (67)$$

We can now show that Theorem 3 is true, which we similarly restate using our subscript-less notation:

**Theorem 11 (Optimistic Exploration)** Let  $\lambda$  satisfy the inequality  $\frac{1}{2} \log(\lambda\pi) \leq \mathbb{H}_\theta[Y|x, a, Z] + D_{\text{KL}}(p_\theta(Z|x, a) \| p_\theta(Z))$ , with  $\pi$  here being the mathematical constant (not the agent’s policy). Then:

$$\mathcal{R}_{\theta, \eta, \nu^*}(x, a) \geq D_{\text{KL}}(\tau(Y|x, a) \| \tau_{\theta, \eta}(Y|x, a)) \quad (68)$$

where  $\tau_{\theta, \eta}(Y|x, a) := \mathbb{E}_{Z \sim p_\theta} p_\eta(Y|x, a, Z)$  denotes the learned environment model. Furthermore, for optimal model parameters  $\theta^*, \eta^*$  we have that the intrinsic reward  $\mathcal{R}_{\theta^*, \eta^*, \nu^*}(x, a) = 0$  for all  $x, a$ .

*Proof.* Use Proposition 8, then the constraint on  $\lambda$ , then Lemma 10:

$$\mathcal{R}_{\theta,\eta,\nu^*}(x, a) := \frac{1}{\lambda} \mathcal{R}_{\theta,\eta}^{\text{rec.}}(x, a) + \lim_{K \rightarrow \infty} \mathcal{R}_{\theta,\nu^*}^{K,\text{inv.}}(x, a) \quad (69)$$

$$= \mathbb{E}_{\substack{Y \sim \tau(\cdot|x,a) \\ Z \sim p_\theta(\cdot|x,a,Y)}} \frac{1}{\lambda} (Y - f_\eta(x, a, Z))^2 + \mathbb{E}_{Z \sim p_\theta(\cdot|x,a)} \text{PMI}_\theta(x, a; Z) \quad (70)$$

$$= \mathbb{E}_{\substack{Y \sim \tau(\cdot|x,a) \\ Z \sim p_\theta(\cdot|x,a,Y)}} \frac{1}{\lambda} (Y - f_\eta(x, a, Z))^2 + D_{\text{KL}}(p_\theta(Z|x, a) \| p_\theta(Z)) \quad (71)$$

$$\geq -\mathbb{E}_{\substack{Y \sim \tau(\cdot|x,a) \\ Z \sim p_\theta(\cdot|x,a,Y)}} \log p_\eta(Y|x, a, Z) - \mathbb{H}_\theta[Y|x, a, Z] \quad (72)$$

$$= -\mathbb{E}_{\substack{Y \sim \tau(\cdot|x,a) \\ Z \sim p_\theta(\cdot|x,a,Y)}} \log p_\eta(Y|x, a, Z) + \mathbb{I}_\theta[Y; Z|x, a] - \mathbb{H}[Y|x, a] \quad (73)$$

$$\begin{aligned} &= -\mathbb{E}_{\substack{Y \sim \tau(\cdot|x,a) \\ Z \sim p_\theta(\cdot|x,a,Y)}} \log p_\eta(Y|x, a, Z) \leftarrow \text{remaining stochasticity} \\ &\quad + \mathbb{E}_{Y \sim \tau(\cdot|x,a)} D_{\text{KL}}(p_\theta(Z|x, a, Y) \| p_\theta(Z|x, a)) \leftarrow \text{hindsight information} \\ &\quad - \mathbb{E}_{Y \sim \tau(\cdot|x,a)} [-\log \tau(Y|x, a)] \leftarrow \text{total stochasticity} \end{aligned} \quad (74)$$

$$\begin{aligned} &\geq -\mathbb{E}_{Y \sim \tau(\cdot|x,a)} [\mathbb{E}_{Z \sim p_\theta(\cdot|x,a,Y)} \log p_\eta(Y|x, a, Z) \\ &\quad - D_{\text{KL}}(p_\theta(Z|x, a, Y) \| p_\theta(Z|x, a)) + D_{\text{KL}}(p_\theta(Z|x, a, Y) \| p_\eta(Z|x, a, Y))] \end{aligned} \quad (75)$$

$$+ \mathbb{E}_{Y \sim \tau(\cdot|x,a)} \log \tau(Y|x, a) \quad (75)$$

$$= -\mathbb{E}_{Y \sim \tau(\cdot|x,a)} \log \mathbb{E}_{Z \sim p_\theta} p_\eta(Y|x, a, Z) + \mathbb{E}_{Y \sim \tau(\cdot|x,a)} \log \tau(Y|x, a) \quad (76)$$

$$= -\mathbb{E}_{Y \sim \tau(\cdot|x,a)} \log \tau_{\theta,\eta}(Y|x, a) + \mathbb{E}_{Y \sim \tau(\cdot|x,a)} \log \tau(Y|x, a) \quad (77)$$

$$= D_{\text{KL}}(\tau(Y|x, a) \| \tau_{\theta,\eta}(Y|x, a)) \quad (78)$$

which completes the proof.  $\square$

The intuition is as follows: Assuming realizability, at convergence “hindsight information” and “total stochasticity” cancel (i.e. neither more nor less), and the “remaining stochasticity” term goes to zero.