## A Appendix

A. 1 Summary of Abbreviation and Notation ..... 16
A. 2 Review of Nonlinear FK Lemma and SB-FBSDE ..... 16
A. 3 Proofs in Main Paper ..... 18
A.3.1 Preliminary ..... 18
A.3.2 Proof of Lemma 1 ..... 19
A.3.3 Proof of Theorem 2 ..... 20
A.3.4 Proof of Proposition 3 ..... 22
A.3.5 Proof of Proposition 4 ..... 22
A. 4 Additional Derivations \& Remarks in Sec. 3 and 4 ..... 24
A.4.1 Hopf-Cole transform ..... 24
A.4.2 Remarks on convergence ..... 24
A.4.3 Functional derivative of MF potential functions ..... 25
A. 5 Experiment Details ..... 25
A.5.1 Setup ..... 25
A.5.2 Additional experiments ..... 27

## A. 1 Summary of Abbreviation and Notation

Table 4: Abbreviation.

| MFGs | Mean-Field Games |
| :--- | :--- |
| SB | Schrödinger Bridge |
| DeepRL | Deep Reinforcement Learning |
| PDEs | Partial Differential Equations |
| HJB | Hamilton-Jacobi-Bellman |
| FP | Fokker-Plank |
| SDEs | Stochastic Differential Equations |
| FBSDEs | Forward-Backward SDEs |
| IPF | Iterative Proportional Fitting |
| MF interaction | Mean-field interaction |
| nonlinear FK | nonlinear Feynman-Kac |
| TD | Temporal Difference |

Table 5: Notation.

| $t$ | time coordinate |
| :--- | :--- |
| $s$ | reversed time coordinate |
| $u(t, x)$ | value function |
| $\rho(t, x)$ | marginal distribution |
| $\rho_{0}, \rho_{\text {target }}$ | initial/target distributions |
| $H$ | Hamiltonian function |
| $F$ | MF interaction function |
| $f$ | MF base drift |
| $\sigma$ | diffusion scaler |
| $(\Psi, \widehat{\Psi})$ | solution to SB PDEs |
| $(Y, Z)$ | nonlinear FK of $\Psi$ |
| $(\widehat{Y}, \widehat{Z})$ | nonlinear FK of $\widehat{\Psi}$ |
| TD | TD target for $Y_{s}$ |
| $\widehat{\mathrm{TD}_{t}}$ | TD target for $\widehat{Y}$ |
| $\theta$ | Parameter of $Y$ (and $Z)$ |
| $\phi$ | Parameter of $\widehat{Y}$ (and $\widehat{Z})$ |

## A. 2 Review of Nonlinear FK Lemma and SB-FBSDE

Lemma 5 (Nonlinear Feynman-Kac Lemma [18, 44, 45]). Let $v \equiv v(x, t)$ be a function that is twice continuously differentiable in $x \in \mathbb{R}^{d}$ and once differentiable in $t \in[0, T]$, i.e., $v \in C^{2,1}\left(\mathbb{R}^{d},[0, T]\right)$.

Consider the following second-order parabolic PDE,

$$
\begin{equation*}
\frac{\partial v}{\partial t}+\frac{1}{2} \operatorname{Tr}\left(\nabla^{2} v G(x, t) G(x, t)^{\top}\right)+\nabla v^{\top} f(x, t)+h\left(x, v, G(x, t)^{\top} \nabla v, t\right)=0, v(T, x)=\varphi(x) \tag{19}
\end{equation*}
$$

where $\nabla^{2}$ denotes the Hessian operator w.r.t. $x$ and the functions $f, G, h$, and $\varphi$ satisfy proper regularity conditions. Specifically, (i) $f, G, h$, and $\varphi$ are continuous, (ii) $f(x, t)$ and $G(x, t)$ are uniformly Lipschitz in $x$, and (iii) $h(x, y, z, t)$ satisfies quadratic growth condition in $z$. Then, (19) exists a unique solution $v$ such that the following stochastic representation (known as the nonlinear Feynman-Kac transformation) holds:

$$
\begin{equation*}
Y_{t}=v\left(X_{t}, t\right), \quad Z_{t}=G\left(X_{t}, t\right)^{\top} \nabla v\left(X_{t}, t\right) \tag{20}
\end{equation*}
$$

where $\left(X_{t}, Y_{t}, Z_{t}\right)$ are the unique adapted solutions to the following FBSDEs:

$$
\begin{align*}
\mathrm{d} X_{t} & =f\left(X_{t}, t\right) \mathrm{d} t+G\left(X_{t}, t\right) \mathrm{d} W_{t}, \quad X_{0}=x_{0} \\
\mathrm{~d} Y_{t} & =-h\left(X_{t}, Y_{t}, Z_{t}, t\right) \mathrm{d} t+Z_{t}^{\top} \mathrm{d} W_{t}, \quad Y_{T}=\varphi\left(X_{T}\right) \tag{21}
\end{align*}
$$

The original deterministic PDE solution $v(x, t)$ can be recovered by taking conditional expectations:

$$
\mathbb{E}\left[Y_{t} \mid X_{t}=x\right]=v(x, t), \quad \mathbb{E}\left[Z_{t} \mid X_{t}=x\right]=G(x, t)^{\top} \nabla v(x, t)
$$

Lemma 5 establishes an intriguing connection between a certain class of (nonlinear) PDEs in (19) and FBSDEs (21) via the nonlinear FK transformation (20). In this work, we adopt a simpler diffusion $G(x, t):=\sigma$ as a time-invariant scalar but note that our derivation can be extended to more general cases straightforwardly.
Viscosity solution. Lemma 5 can be extended to viscosity solutions when the classical solution does not exist. In which case, we will have $v(x, t)=\lim _{\epsilon \rightarrow \infty} v^{\epsilon}(x, t)$ converge uniformly in $(x, t)$ over a compact set, where $v^{\epsilon}(x, t)$ is the classical solution to (19) with $\left(f_{\epsilon}, G_{\epsilon}, h_{\epsilon}, \varphi_{\epsilon}\right)$ converge uniformly toward $(f, G, h, \varphi)$ over the compact set; see $[18,59,60]$ for a complete discussion.
SB-FBSDE [27]. SB-FBSDE is a new class of generative models that, inspiring by the recent advance of understanding deep learning through the optimal control perspective [61-63], adopts Lemma 5 to generalize the score-based diffusion models. Since the PDEs ( $\frac{\partial \Psi}{\partial t}, \frac{\partial \widehat{\Psi}}{\partial t}$ ) appearing in the vanilla SB (4) are both of the parabolic form (19), one can apply Lemma 5 and derive the corresponding nonlinear generators $h$. This, as shown in SB-FBSDE [27], leads to the following FBSDEs:

$$
\left\{\begin{align*}
\mathrm{d} X_{t} & =\left(f_{t}+\sigma Z_{t}\right) \mathrm{d} t+\sigma \mathrm{d} W_{t}  \tag{22a}\\
\mathrm{~d} Y_{t} & =\frac{1}{2}\left\|Z_{t}\right\|^{2} \mathrm{~d} t+Z_{t}^{\top} \mathrm{d} W_{t} \\
\mathrm{~d} \widehat{Y}_{t} & =\left(\frac{1}{2}\left\|\widehat{Z}_{t}\right\|^{2}+\nabla \cdot\left(\sigma \widehat{Z}_{t}-f_{t}\right)+\widehat{Z}_{t}^{\top} Z_{t}\right) \mathrm{d} t+\widehat{Z}_{t}^{\top} \mathrm{d} W_{t}
\end{align*}\right.
$$

Further, the nonlinear FK transformation reads

$$
\begin{array}{ll}
Y_{t}=\log \Psi\left(X_{t}, t\right), & Z_{t}=\sigma \nabla \log \Psi\left(X_{t}, t\right), \\
\widehat{Y}_{t}=\log \widehat{\Psi}\left(X_{t}, t\right), & \widehat{Z}_{t}=\sigma \nabla \log \widehat{\Psi}\left(X_{t}, t\right),
\end{array}
$$

which immediately suggests that

$$
\begin{equation*}
\mathbb{E}\left[Y_{t} \mid X_{t}=x\right]=\log \Psi(x, t), \quad \mathbb{E}\left[\widehat{Y}_{t} \mid X_{t}=x\right]=\log \widehat{\Psi}(x, t) \tag{23}
\end{equation*}
$$

It can be readily seen that (22) is a special case of our Theorem 2 when the MF interaction $F(x, \rho)$, which plays a crucial role in MFGs, vanishes. Since SB-FBSDE was primarily developed in the context of generative modeling [64], its training relies on computing the log-likelihood at the boundaries. These log-likelihoods can be obtained by noticing that $\log \rho(x, t)=\mathbb{E}\left[Y_{t}+\widehat{Y}_{t} \mid X_{t}=x\right]$, as implied by (23) and (8). When $\widehat{Z}_{\phi}\left(X_{t}, t\right) \approx \widehat{Z}_{t}$ and $Z_{\theta}\left(\bar{X}_{s}, s\right) \approx Z_{s}$, the training objectives of SB-FBSDE can be computed as the parametrized variational lower-bounds:

$$
\begin{align*}
& \log \rho_{0}(\phi ; x) \geq \mathcal{L}_{\mathrm{IPF}}(\phi):=\mathbb{E}\left[Y_{t}^{\theta}+\widehat{Y}_{t}^{\phi} \mid X_{t}=x, t=0\right]=\int_{t} \mathbb{E}\left[\mathrm{~d} Y_{t}^{\theta}+\mathrm{d} \widehat{Y}_{t}^{\phi} \mid X_{0}=x\right]  \tag{24a}\\
& \log \rho_{T}(\theta ; x) \geq \mathcal{L}_{\mathrm{IPF}}(\theta):=\mathbb{E}\left[Y_{s}^{\theta}+\widehat{Y}_{s}^{\phi} \mid \bar{X}_{s}=x, s=0\right]=\int_{s} \mathbb{E}\left[\mathrm{~d} Y_{s}^{\theta}+\mathrm{d} \widehat{Y}_{s}^{\phi} \mid \bar{X}_{0}=x\right] \tag{24b}
\end{align*}
$$

Invoking (22) to expand the r.h.s. of (24) leads to the expression in (6):

$$
\begin{aligned}
& \mathcal{L}_{\mathrm{IPF}}(\theta)=\int_{0}^{T} \mathbb{E}_{(3 \mathrm{~b})}\left[\frac{1}{2}\left\|Z_{\theta}\left(\bar{X}_{s}, s\right)\right\|_{2}^{2}+Z_{\theta}\left(\bar{X}_{s}, s\right)^{\top} \widehat{Z}_{\phi}\left(\bar{X}_{s}, s\right)+\nabla \cdot\left(\sigma Z_{\theta}\left(\bar{X}_{s}, s\right)+f\right)\right] \mathrm{d} s \\
& \mathcal{L}_{\mathrm{IPF}}(\phi)=\int_{0}^{T} \mathbb{E}_{(3 \mathrm{aa})}\left[\frac{1}{2}\left\|\widehat{Z}_{\phi}\left(X_{t}, t\right)\right\|_{2}^{2}+\widehat{Z}_{\phi}\left(X_{t}, t\right)^{\top} Z_{\theta}\left(X_{t}, t\right)+\nabla \cdot\left(\sigma \widehat{Z}_{\phi}\left(X_{t}, t\right)-f\right)\right] \mathrm{d} t
\end{aligned}
$$

Since (24) concern only the integration over the expectations, i.e., $\int \mathbb{E}[\mathrm{d} Y+\mathrm{d} \widehat{Y}]$, the solutions $\left(Y_{t}, \widehat{Y}_{t}\right)$ to the SDEs (22b, 22c) were never computed explicitly in SB-FBSDE, This is in contrast to our DeepGSB, which, crucially, requires computing $\left(Y_{t}, \widehat{Y}_{t}\right)$ explicitly and regress their values with TD objectives, so that the stochastic dynamics of $\mathrm{d} Y$ and $\mathrm{d} \widehat{Y}$ are respectively respected.

## A. 3 Proofs in Main Paper

Throughout this section, we will denote the parameterized forward and backward SDEs by

$$
\begin{align*}
& \mathrm{d} X_{t}^{\theta}=\left(f_{t}+\sigma Z_{\theta}\left(X_{t}^{\theta}, t\right)\right) \mathrm{d} t+\sigma \mathrm{d} W_{t}  \tag{25a}\\
& \mathrm{~d} \bar{X}_{s}^{\phi}=\left(-f_{s}+\sigma \widehat{Z}_{\phi}\left(\bar{X}_{s}^{\phi}, t\right)\right) \mathrm{d} s+\sigma \mathrm{d} W_{s} \tag{25b}
\end{align*}
$$

and denote their time-marginal densities respectively as $q^{\theta}$ and $q^{\phi}$.

## A.3.1 Preliminary

We first restate some useful lemmas that will appear in the proceeding proofs.
Lemma 6 (Itô formula [46]). Let $X_{t}$ be the solution to the Itô SDE:

$$
\mathrm{d} X_{t}=f\left(X_{t}, t\right) \mathrm{d} t+\sigma\left(X_{t}, t\right) \mathrm{d} W_{t}
$$

Then, the stochastic process $v\left(X_{t}, t\right)$, where $v \in C^{2,1}\left(\mathbb{R}^{d},[0, T]\right)$, is also an Itô process satisfying

$$
\begin{equation*}
\mathrm{d} v\left(X_{t}, t\right)=\frac{\partial v\left(X_{t}, t\right)}{\partial t} \mathrm{~d} t+\left[\nabla v\left(X_{t}, t\right)^{\top} f+\frac{1}{2} \operatorname{Tr}\left[\sigma^{\top} \nabla^{2} v\left(X_{t}, t\right) \sigma\right]\right] \mathrm{d} t+\left[\nabla v\left(X_{t}, t\right)^{\top} \sigma\right] \mathrm{d} W_{t} . \tag{26}
\end{equation*}
$$

Lemma 7. The following equality holds at any point $x \in \mathbb{R}^{n}$ such that $p(x) \neq 0$.

$$
\frac{1}{p(x)} \Delta p(x)=\|\nabla \log p(x)\|^{2}+\Delta \log p(x)
$$

Proof. $\frac{1}{p(x)} \Delta p(x)=\frac{1}{p(x)} \nabla \cdot \nabla p(x)=\frac{1}{p(x)} \nabla \cdot(p(x) \nabla \log p(x))$. Applying chain rule to the divergence yields the desired result.
Lemma 8 (Vargas [38], Proposition 1, Sec 6.3.1).
$\mathrm{d} \log q_{t}^{\phi}=\left[\nabla \cdot\left(\sigma \widehat{Z}_{\phi}-f_{t}\right)+\sigma\left(Z_{\theta}+\widehat{Z}_{\phi}\right)^{\top} \nabla \log q_{t}^{\phi}-\frac{1}{2}\left\|\sigma \nabla \log q_{t}^{\phi}\right\|^{2}\right] \mathrm{d} t+\sigma \nabla \log q_{t}^{\phi^{\top}} \mathrm{d} W_{t}$.
Proof. Invoking Ito lemma w.r.t. the parameterized forward $\operatorname{SDE}$ (25a),

$$
\mathrm{d} \log q_{t}^{\phi}=\left[\frac{\partial \log q_{t}^{\phi}}{\partial t}+\nabla \log q_{t}^{\phi^{\top}}\left(f_{t}+\sigma Z_{\theta}\right)+\frac{\sigma^{2}}{2} \Delta \log q_{t}^{\phi}\right] \mathrm{d} t+\sigma \nabla \log q_{t}^{\phi^{\top}} \mathrm{d} W_{t}
$$

where $\frac{\partial \log q_{t}^{\phi}}{\partial t}$ obeys (see Eq 13.4 in Nelson [65]):

$$
\begin{aligned}
-\frac{\partial q_{t}^{\phi}}{\partial t} & =-\nabla \cdot\left(\left(\sigma \widehat{Z}_{\phi}-f_{t}\right) q_{t}^{\phi}\right)+\frac{\sigma^{2}}{2} \Delta q_{t}^{\phi} \\
\Rightarrow \frac{\partial \log q_{t}^{\phi}}{\partial t} & =\nabla \cdot\left(\sigma \widehat{Z}_{\phi}-f_{t}\right)+\left(\sigma \widehat{Z}_{\phi}-f_{t}\right)^{\top} \nabla \log q_{t}^{\phi}-\frac{\sigma^{2} \Delta q_{t}^{\phi}}{2 q_{t}^{\phi}}
\end{aligned}
$$

Substituting the above relation yields the desired results.

Proposition 9 (Vargas [38], Proposition 1 in Sec 6.3.1).

$$
D_{\mathrm{KL}}\left(q^{\theta} \| q^{\phi}\right)=\int_{0}^{T} \mathbb{E}_{q_{t}^{\theta}}\left[\frac{1}{2}\left\|\widehat{Z}_{\phi}+Z_{\theta}\right\|^{2}+\nabla \cdot\left(\sigma \widehat{Z}_{\phi}-f_{t}\right)\right] \mathrm{d} t+\mathbb{E}_{q_{0}^{\theta}}\left[\log \rho_{0}\right]-\mathbb{E}_{q_{T}^{\theta}}\left[\log \rho_{\text {target }}\right]
$$

Proof. Recall that the parametrized backward $\operatorname{SDE}(25 b)$ can be reversed $[64,66]$ as

$$
\mathrm{d} \bar{X}_{t}^{\phi}=\left(f_{t}-\sigma \widehat{Z}_{\phi}\left(\bar{X}_{t}^{\phi}, t\right)+\sigma^{2} \nabla \log q^{\phi}\left(\bar{X}_{t}^{\phi}, t\right)\right) \mathrm{d} t+\sigma \mathrm{d} W_{t} .
$$

Then, we have

$$
\begin{aligned}
& D_{\mathrm{KL}}\left(q^{\theta} \| q^{\phi}\right) \\
= & \int_{0}^{T} \mathbb{E}_{q_{t}^{\theta}}\left[\frac{1}{2}\left\|\widehat{Z}_{\phi}+Z_{\theta}-\sigma \nabla \log q_{t}^{\phi}\right\|^{2}\right] \mathrm{d} t+D_{\mathrm{KL}}\left(\rho_{0} \| q_{t=0}^{\phi}\right) \\
= & \int_{0}^{T} \mathbb{E}_{q_{t}^{\theta}}\left[\frac{1}{2}\left\|\widehat{Z}_{\phi}+Z_{\theta}\right\|^{2}-\sigma\left(\widehat{Z}_{\phi}+Z_{\theta}\right)^{T} \nabla \log q_{t}^{\phi}+\frac{1}{2}\left\|\sigma \nabla \log q_{t}^{\phi}\right\|^{2}\right] \mathrm{d} t+D_{\mathrm{KL}}\left(\rho_{0} \| q_{t=0}^{\phi}\right) \\
\stackrel{(*)}{=} & \int_{0}^{T} \mathbb{E}_{q_{t}^{\theta}}\left[\frac{1}{2}\left\|\widehat{Z}_{\phi}+Z_{\theta}\right\|^{2}+\nabla \cdot\left(\sigma \widehat{Z}_{\phi}-f_{t}\right)\right] \mathrm{d} t-\mathbb{E}_{q^{\theta}}\left[\int_{0}^{T} \mathrm{~d} \log q_{t}^{\phi}\right]+D_{\mathrm{KL}}\left(\rho_{0} \| q_{t=0}^{\phi}\right) \\
= & \int_{0}^{T} \mathbb{E}_{q_{t}^{\theta}}\left[\frac{1}{2}\left\|\widehat{Z}_{\phi}+Z_{\theta}\right\|^{2}+\nabla \cdot\left(\sigma \widehat{Z}_{\phi}-f_{t}\right)\right] \mathrm{d} t+\mathbb{E}_{q_{0}^{\theta}}\left[\log \rho_{0}\right]-\mathbb{E}_{q_{T}^{\theta}}\left[\log \rho_{\text {target }}\right],
\end{aligned}
$$

where $(*)$ is due to Lemma 8.

## A.3.2 Proof of Lemma 1

Proof. Substituting $\mathcal{L}_{\mathrm{IPF}}(\phi)$ into Proposition 9 and dropping all terms independent of $\phi$ readily yields $D_{\mathrm{KL}}\left(q^{\theta} \| q^{\phi}\right) \propto \mathcal{L}_{\mathrm{IPF}}(\phi)$. A similar relation can be derived between $D_{\mathrm{KL}}\left(q^{\phi} \| q^{\theta}\right)$.

Remark (an alternative simpler proof). Suppose $\left(Z_{\theta}, q^{\theta}\right)$ and $\left(\widehat{Z}_{\phi}, q^{\phi}\right)$ satisfy proper regularity such that $\forall t, s \in[0, T], \quad \exists k>0: q^{\theta}(x, t)=\mathcal{O}\left(\exp ^{-\|x\|_{k}^{2}}\right), q^{\phi}(x, s)=\mathcal{O}\left(\exp ^{-\|x\|_{k}^{2}}\right)$ as $x \rightarrow \infty$. Then, an alternative proof using integration by part goes as follows: Recall that the parametrized forward SDE in (25a) can be reversed [64, 66] as

$$
\mathrm{d} X_{s}^{\theta}=\left(-f_{s}-\sigma Z_{\theta}\left(X_{s}^{\theta}, s\right)+\sigma^{2} \nabla \log q^{\theta}\left(X_{s}^{\theta}, s\right)\right) \mathrm{d} s+\sigma \mathrm{d} W_{s}
$$

Then, the KL divergence can be computed as

$$
\begin{aligned}
& D_{\mathrm{KL}}\left(q^{\theta} \| q^{\phi}\right) \\
& \stackrel{(*)}{=} \mathbb{E}_{q^{\theta}}\left[\int_{0}^{T} \frac{1}{2 \sigma^{2}}\left\|\sigma \widehat{Z}_{\phi}+\sigma Z_{\theta}-\sigma^{2} \nabla \log q_{s}^{\theta}\right\|^{2} \mathrm{~d} s\right]+D_{\mathrm{KL}}\left(q_{s=0}^{\theta} \| \rho_{\text {target }}\right) \\
&= \int_{0}^{T} \mathbb{E}_{q_{s}^{\theta}}\left[\frac{1}{2}\left\|\widehat{Z}_{\phi}+Z_{\theta}\right\|^{2}-\sigma\left(\widehat{Z}_{\phi}+Z_{\theta}\right)^{\top} \nabla \log q_{s}^{\theta}+\frac{1}{2}\left\|\sigma \nabla \log q_{s}^{\theta}\right\|^{2}\right] \mathrm{d} s+D_{\mathrm{KL}}\left(q_{0}^{\theta} \| \rho_{\mathrm{target}}\right) \\
&= \int_{0}^{T} \mathbb{E}_{q_{s}^{\theta}}\left[\frac{1}{2}\left\|\widehat{Z}_{\phi}\right\|^{2}+\widehat{Z}_{\phi}^{\top} Z_{\theta}-\sigma \widehat{Z}_{\phi}^{\top} \nabla \log q_{s}^{\theta}\right] \mathrm{d} s+\mathcal{O}(1) \\
& \stackrel{(* *)}{=} \int_{0}^{T} \mathbb{E}_{q_{s}^{\theta}}\left[\frac{1}{2}\left\|\widehat{Z}_{\phi}\right\|^{2}+\widehat{Z}_{\phi}^{\top} Z_{\theta}+\sigma \nabla \cdot \widehat{Z}_{\phi}\right] \mathrm{d} s+\mathcal{O}(1), \\
& \propto \mathcal{L}_{\mathrm{IPF}}(\phi)
\end{aligned}
$$

where $\left({ }^{*}\right)$ is due to the Girsanov's Theorem [67] and $\left({ }^{* *}\right)$ is due to integration by parts. $\mathcal{O}(1)$ collects terms independent of $\phi$. Notice that the boundary terms vanish due to the additional regularity assumptions on $q^{\theta}$ and $q^{\phi}$. Similar transformations have been adopted in e.g., Theorem 1 in Song et al. [68] or Theorem 3 in Huang et al. [69].

## A.3.3 Proof of Theorem 2

Proof. Apply the Itô formula to $v:=\log \Psi\left(X_{t}, t\right)$, where $X_{t}$ follows (3a),

$$
\mathrm{d} \log \Psi=\frac{\partial \log \Psi}{\partial t} \mathrm{~d} t+\left[\nabla \log \Psi^{\top}\left(f+\sigma^{2} \nabla \log \Psi\right)+\frac{\sigma^{2}}{2} \Delta \log \Psi\right] \mathrm{d} t+\sigma \nabla \log \Psi^{\top} \mathrm{d} W_{t}
$$

and notice that the PDE of $\frac{\partial \log \Psi}{\partial t}$ obeys

$$
\frac{\partial \log \Psi}{\partial t}=\frac{1}{\Psi}\left(-\nabla \Psi^{\top} f-\frac{\sigma^{2}}{2} \Delta \Psi+F \Psi\right)=-\nabla \log \Psi^{\top} f-\frac{\sigma^{2}}{2}\|\nabla \log \Psi\|^{2}-\frac{\sigma^{2}}{2} \Delta \log \Psi+F .
$$

This yields

$$
\begin{equation*}
\mathrm{d} \log \Psi=\left[\frac{1}{2}\|\sigma \nabla \log \Psi\|^{2}+F\right] \mathrm{d} t+\sigma \nabla \log \Psi^{\top} \mathrm{d} W_{t} \tag{28}
\end{equation*}
$$

Now, apply the same Itô formula by instead substituting $v:=\log \widehat{\Psi}\left(X_{t}, t\right)$, where $X_{t}$ follows (3a),

$$
\mathrm{d} \log \widehat{\Psi}=\frac{\partial \log \widehat{\Psi}}{\partial t} \mathrm{~d} t+\left[\nabla \log \widehat{\Psi}^{\top}\left(f+\sigma^{2} \nabla \log \Psi\right)+\frac{\sigma^{2}}{2} \Delta \log \widehat{\Psi}\right] \mathrm{d} t+\sigma \nabla \log \widehat{\Psi}^{\top} \mathrm{d} W_{t}
$$

and notice that the PDE of $\frac{\partial \log \widehat{\Psi}}{\partial t}$ obeys

$$
\begin{aligned}
\frac{\partial \log \widehat{\Psi}}{\partial t} & =\frac{1}{\widehat{\Psi}}\left(-\nabla \cdot(\widehat{\Psi} f)+\frac{\sigma^{2}}{2} \Delta \widehat{\Psi}-F \widehat{\Psi}\right) \\
& =-\nabla \log \widehat{\Psi}^{\top} f-\nabla \cdot f+\frac{\sigma^{2}}{2}\|\nabla \log \widehat{\Psi}\|^{2}+\frac{\sigma^{2}}{2} \Delta \log \widehat{\Psi}-F
\end{aligned}
$$

This yields

$$
\begin{align*}
\mathrm{d} \log \widehat{\Psi} & =\left[-\nabla \cdot f+\frac{\sigma^{2}}{2}\|\nabla \log \widehat{\Psi}\|^{2}+\sigma^{2} \nabla \log \widehat{\Psi}^{\top} \nabla \log \Psi+\sigma^{2} \Delta \log \widehat{\Psi}-F\right] \mathrm{d} t+\sigma \nabla \log \widehat{\Psi}^{\top} \mathrm{d} W_{t} \\
& =\left[\nabla \cdot\left(\sigma^{2} \nabla \log \widehat{\Psi}-f\right)+\frac{\sigma^{2}}{2}\|\nabla \log \widehat{\Psi}\|^{2}+\sigma^{2} \nabla \log \widehat{\Psi}^{\top} \nabla \log \Psi-F\right] \mathrm{d} t+\sigma \nabla \log \widehat{\Psi}^{\top} \mathrm{d} W_{t} . \tag{29}
\end{align*}
$$

Finally, with the nonlinear FK transformation in (10), i.e.,

$$
\begin{array}{ll}
Y_{t} \equiv Y\left(X_{t}, t\right)=\log \Psi\left(X_{t}, t\right), & Z_{t} \equiv Z\left(X_{t}, t\right)=\sigma \nabla \log \Psi\left(X_{t}, t\right), \\
\widehat{Y}_{t} \equiv \widehat{Y}\left(X_{t}, t\right)=\log \widehat{\Psi}\left(X_{t}, t\right), & \widehat{Z}_{t} \equiv \widehat{Z}\left(X_{t}, t\right)=\sigma \nabla \log \widehat{\Psi}\left(X_{t}, t\right),
\end{array}
$$

we can rewrite (3a, 28, 29) as the FBSDEs system in (11).

$$
\begin{aligned}
\mathrm{d} X_{t} & =\left(f_{t}+\sigma Z_{t}\right) \mathrm{d} t+\sigma \mathrm{d} W_{t} \\
\mathrm{~d} Y_{t} & =\left[\frac{1}{2}\left\|Z_{t}\right\|^{2}+F_{t}\right] \mathrm{d} t+Z_{t}^{\top} \mathrm{d} W_{t} \\
\mathrm{~d} \widehat{Y}_{t} & =\left[\frac{1}{2}\left\|\widehat{Z}_{t}\right\|^{2}+\widehat{Z}_{t}^{\top} Z_{t}+\nabla \cdot\left(\sigma \widehat{Z}_{t}-f_{t}\right)-F_{t}\right]+\widehat{Z}^{\top} \mathrm{d} W_{t}
\end{aligned}
$$

where

$$
f_{t}:=f\left(X_{t}, \exp \left(Y_{t}+\widehat{Y}_{t}\right)\right), \quad F_{t}:=F\left(X_{t}, \exp \left(Y_{t}+\widehat{Y}_{t}\right)\right)
$$

Derivation of the second FBSDEs system in (12) follows a similar flow, except that we need to rebase the PDEs (9) to the "reversed" time coordinate $s:=T-t$. This can be done by reformulating the HJB and FP PDEs in (7) under the $s$ coordinate, then applying the following Hopf-Cole transform:

$$
\begin{equation*}
\widehat{\Psi}(x, s):=\exp (-u(x, s)), \quad \Psi(x, s):=\rho(x, s) \exp (u(x, s)) \tag{31}
\end{equation*}
$$

Notice that we flip the role of $\widehat{\Psi}(x, s)$ and $\Psi(x, s)$ as the former now relates to the policy appearing in (3b). Omitting the computation similar to Appendix A.4.1, we arrive at the following:

$$
\left\{\begin{array}{l}
\frac{\partial \widehat{\Psi}(x, s)}{\partial s}=\nabla \widehat{\Psi}^{\top} f-\frac{1}{2} \sigma^{2} \Delta \widehat{\Psi}+F \widehat{\Psi}  \tag{32}\\
\frac{\partial \Psi(x, s)}{\partial s}=\nabla \cdot(\Psi f)+\frac{1}{2} \sigma^{2} \Delta \Psi-F \Psi
\end{array} \quad \text { s.t. } \begin{array}{c}
\widehat{\Psi}(\cdot, 0) \Psi(\cdot, 0)=\rho_{\text {target }} \\
\widehat{\Psi}(\cdot, T) \Psi(\cdot, T)=\rho_{0}
\end{array}\right.
$$

Apply the Itô formula to $v:=\log \Psi\left(\bar{X}_{s}, s\right)$, where $\bar{X}_{s}$ evolves along the reversed $\operatorname{SDE}$ (3b).

$$
\mathrm{d} \log \Psi=\frac{\partial \log \Psi}{\partial s} \mathrm{~d} s+\left[\nabla \log \Psi^{\top}\left(-f+\sigma^{2} \nabla \log \widehat{\Psi}\right)+\frac{\sigma^{2}}{2} \Delta \log \Psi\right] \mathrm{d} s+\sigma \nabla \log \Psi^{\top} \mathrm{d} W_{s}
$$

and notice that the PDE of $\frac{\partial \log \Psi}{\partial s}$ now obeys

$$
\begin{aligned}
\frac{\partial \log \Psi}{\partial s} & =\frac{1}{\Psi}\left(\nabla \cdot(\Psi f)+\frac{\sigma^{2}}{2} \Delta \Psi-F \Psi\right) \\
& =\nabla \log \Psi^{\top} f+\nabla \cdot f+\frac{\sigma^{2}}{2}\|\nabla \log \Psi\|^{2}+\frac{\sigma^{2}}{2} \Delta \log \Psi-F .
\end{aligned}
$$

This yields

$$
\begin{align*}
\mathrm{d} \log \Psi & =\left[\nabla \cdot f+\frac{\sigma^{2}}{2}\|\nabla \log \Psi\|^{2}+\sigma^{2} \nabla \log \Psi^{\top} \nabla \log \widehat{\Psi}+\sigma^{2} \Delta \log \Psi-F\right] \mathrm{d} s+\sigma \nabla \log \Psi^{\top} \mathrm{d} W_{s} \\
& =\left[\nabla \cdot\left(f+\sigma^{2} \nabla \log \Psi\right)+\frac{\sigma^{2}}{2}\|\nabla \log \Psi\|^{2}+\sigma^{2} \nabla \log \Psi^{\top} \nabla \log \widehat{\Psi}-F\right] \mathrm{d} s+\sigma \nabla \log \Psi^{\top} \mathrm{d} W_{s} \tag{33}
\end{align*}
$$

Similarly, apply the Itô formula to $v:=\log \widehat{\Psi}\left(\bar{X}_{s}, s\right)$, where $\bar{X}_{s}$ follows the same reversed $\operatorname{SDE}$ (3b).

$$
\mathrm{d} \log \widehat{\Psi}=\frac{\partial \log \widehat{\Psi}}{\partial s} \mathrm{~d} s+\left[\nabla \log \widehat{\Psi}^{\top}\left(-f+\sigma^{2} \nabla \log \widehat{\Psi}\right)+\frac{\sigma^{2}}{2} \Delta \log \widehat{\Psi}\right] \mathrm{d} s+\sigma \nabla \log \widehat{\Psi}^{\top} \mathrm{d} W_{s}
$$

and notice that the PDE of $\frac{\partial \log \widehat{\Psi}}{\partial s}$ obeys

$$
\frac{\partial \log \widehat{\Psi}}{\partial s}=\frac{1}{\widehat{\Psi}}\left(\nabla \widehat{\Psi}^{\top} f-\frac{\sigma^{2}}{2} \Delta \widehat{\Psi}+F \widehat{\Psi}\right)=\nabla \log \widehat{\Psi}^{\top} f-\frac{\sigma^{2}}{2}\|\nabla \log \widehat{\Psi}\|^{2}-\frac{\sigma^{2}}{2} \Delta \log \widehat{\Psi}+F
$$

This yields

$$
\begin{equation*}
\mathrm{d} \log \widehat{\Psi}=\left[\frac{1}{2}\|\sigma \nabla \log \widehat{\Psi}\|^{2}+F\right] \mathrm{d} s+\sigma \nabla \log \widehat{\Psi}^{\top} \mathrm{d} W_{s} \tag{34}
\end{equation*}
$$

Finally, with a nonlinear FK transformation similar to (10),

$$
\begin{array}{ll}
Y_{s} \equiv Y\left(\bar{X}_{s}, s\right)=\log \Psi\left(\bar{X}_{s}, s\right), & Z_{s} \equiv Z\left(\bar{X}_{s}, s\right)=\sigma \nabla \log \Psi\left(\bar{X}_{s}, s\right), \\
\widehat{Y}_{s} \equiv \widehat{Y}\left(\bar{X}_{s}, s\right)=\log \widehat{\Psi}\left(\bar{X}_{s}, s\right), & \widehat{Z}_{s} \equiv \widehat{Z}\left(\bar{X}_{s}, s\right)=\sigma \nabla \log \widehat{\Psi}\left(\bar{X}_{s}, s\right), \tag{35}
\end{array}
$$

we can rewrite $(3 \mathrm{~b}, 33,34)$ as the second FBSDEs system in (12).

$$
\begin{aligned}
\mathrm{d} \bar{X}_{s} & =\left(-f_{s}+\sigma \widehat{Z}_{s}\right) \mathrm{d} s+\sigma \mathrm{d} W_{s} \\
\mathrm{~d} Y_{s} & =\left(\frac{1}{2}\left\|Z_{s}\right\|^{2}+\nabla \cdot\left(\sigma Z_{s}+f_{s}\right)+Z_{s}^{\top} \widehat{Z}_{s}-F_{s}\right) \mathrm{d} s+Z_{s}^{\top} \mathrm{d} W_{s} \\
\mathrm{~d} \widehat{Y}_{s} & =\left(\frac{1}{2}\left\|\widehat{Z}_{s}\right\|^{2}+F_{s}\right) \mathrm{d} s+\widehat{Z}_{s}^{\top} \mathrm{d} W_{s}
\end{aligned}
$$

where

$$
f_{s}:=f\left(\bar{X}_{s}, \exp \left(Y_{s}+\widehat{Y}_{s}\right)\right), \quad F_{s}:=F\left(\bar{X}_{s}, \exp \left(Y_{s}+\widehat{Y}_{s}\right)\right)
$$

We conclude the proof.

## A.3.4 Proof of Proposition 3

Proof. We will only prove the TD objective (14a) for the time coordinate $t$, as all derivations can be adopted similarly to its reversed coordinate $s:=T-t$.

Given a realization of the parametrized $\operatorname{SDE}$ (11a) w.r.t. some fixed step size $\delta t$, i.e.,

$$
X_{t+\delta t}^{\theta}=X_{t}^{\theta}+\left(f_{t}+\sigma Z_{\theta}\left(X_{t}^{\theta}, t\right)\right) \delta t+\delta W_{t}, \quad \delta W_{t} \sim \mathcal{N}(\mathbf{0}, \delta t \boldsymbol{I})
$$

we can represent the trajectory compactly by a sequence of tuples $\boldsymbol{X}_{t}^{\theta} \equiv\left(X_{t}^{\theta}, Z_{t}^{\theta}, \delta W_{t}\right)$ sampled on some discrete time grids, $t \in\{0, \delta t, \cdots, T-\delta t, T\}$. The incremental change of $\widehat{Y}_{t}$, i.e., the r.h.s. of (11c), can then be computed by
$\delta \widehat{Y}_{t}\left(\boldsymbol{X}_{t}^{\theta}\right):=\left(\frac{1}{2}\left\|\widehat{Z}\left(X_{t}^{\theta}, t\right)\right\|^{2}+\nabla \cdot\left(\sigma \widehat{Z}\left(X_{t}^{\theta}, t\right)-f_{t}\right)+\widehat{Z}\left(X_{t}^{\theta}, t\right)^{\top} Z_{t}^{\theta}-F_{t}\right) \delta t+\widehat{Z}\left(X_{t}^{\theta}, t\right)^{\top} \delta W_{t}$,
where $\widehat{Z}(\cdot, \cdot)$ is the (parametrized) backward policy and we denote $Z_{t}^{\theta}:=Z_{\theta}\left(X_{t}^{\theta}, t\right)$ for simplicity. At the equilibrium when the FBSDE system (11) is satisfied, the SDE (11c) must hold. This suggests the following equality:

$$
\begin{equation*}
\widehat{Y}\left(X_{t+\delta t}^{\theta}, t+\delta t\right)=\widehat{Y}\left(X_{t}^{\theta}, t\right)+\delta \widehat{Y}_{t}\left(\boldsymbol{X}_{t}^{\theta}\right) \tag{37}
\end{equation*}
$$

Hence, we can interpret the r.h.s. of (37) as the single-step TD target $\widehat{\mathrm{TD}}_{t+\delta t}^{\text {single }}$, which yields the expression in (14a). The multi-step TD target can be constructed accordingly as standard practices [51,52], and either TD target can be used to construct the TD objective for the parametrized function $\widehat{Y}_{\phi} \approx \widehat{Y}$, which further yields (16).

## A.3.5 Proof of Proposition 4

Proof. We first prove the necessity. Suppose the parametrized functions $\left(Y_{\theta}, Z_{\theta}, \widehat{Y}_{\phi}, \widehat{Z}_{\phi}\right)$ satisfy the SDEs in $(11,12)$, it can be readily seen that the TD objectives $\mathcal{L}_{\mathrm{TD}}(\phi)$ and $\mathcal{L}_{\mathrm{TD}}(\theta)$ shall both be minimized, as the parametrized functions satisfy (11c,12b). Next, notice that (11) implies

$$
\begin{aligned}
Y_{T}^{\theta}+\widehat{Y}_{T}^{\phi} & =\left(Y_{0}^{\theta}+\int_{0}^{T} \mathrm{~d} Y_{t}^{\theta}\right)+\left(\widehat{Y}_{0}^{\phi}+\int_{0}^{T} \mathrm{~d} \widehat{Y}_{t}^{\phi}\right) \\
\Rightarrow 0 & =\mathbb{E}_{q^{\theta}}\left[\left(Y_{0}^{\theta}+\widehat{Y}_{0}^{\phi}\right)+\int_{0}^{T}\left(\mathrm{~d} Y_{t}^{\theta}+\mathrm{d} \widehat{Y}_{t}^{\phi}\right)-\left(Y_{T}^{\theta}+\widehat{Y}_{T}^{\phi}\right)\right] \\
& \stackrel{(*)}{=} \mathbb{E}_{q_{0}^{\theta}}\left[\log \rho_{0}\right]+\int_{0}^{T} \mathbb{E}_{q_{t}^{\theta}}\left[\frac{1}{2}\left\|Z_{t}^{\theta}+\widehat{Z}_{t}^{\phi}\right\|^{2}+\nabla \cdot\left(\sigma \widehat{Z}_{t}^{\phi}-f_{t}\right)\right] \mathrm{d} t-\mathbb{E}_{q_{T}^{\theta}}\left[\log \rho_{\text {target }}\right] \\
& \stackrel{(* *)}{=} \mathbb{E}_{q_{0}^{\theta}}\left[\log \rho_{0}\right]+D_{\mathrm{KL}}\left(q^{\theta} \| q^{\phi}\right)-\mathbb{E}_{q^{\theta}}\left[\log \frac{\rho_{0}}{\rho_{\text {target }}}\right]-\mathbb{E}_{q_{T}^{\theta}}\left[\log \rho_{\text {target }}\right] \\
& =D_{\mathrm{KL}}\left(q^{\theta} \| q^{\phi}\right),
\end{aligned}
$$

where $\left({ }^{*}\right)$ is due to (11b,11c) and (**) invokes Proposition 9. The fact that $\mathcal{L}_{\mathrm{IPF}}(\phi) \propto D_{\mathrm{KL}}\left(q^{\theta} \| q^{\phi}\right)=$ 0 (recall Lemma 1) suggests that the objective $\mathcal{L}_{\text {IPF }}(\phi)$ is minimized when (11) holds. Finally, as similar arguments can be adopted to $\mathcal{L}_{\mathrm{IPF}}(\theta) \propto D_{\mathrm{KL}}\left(q^{\phi} \| q^{\theta}\right)=0$ when (12) holds, we conclude that all losses are minimized when the parameterized functions satisfy the $\operatorname{FBSDE}$ systems $(11,12)$.

We proceed to proving the sufficiency, which is more involved. First, notice that

$$
\begin{align*}
& \mathcal{L}_{\mathrm{IPF}}(\phi) \text { is minimized } \Leftrightarrow D_{\mathrm{KL}}\left(q^{\theta} \| q^{\phi}\right)=0 \Leftrightarrow \forall s \in[0, T], Z_{s}^{\theta}+\widehat{Z}_{s}^{\phi}-\sigma \nabla \log q_{s}^{\theta}=0,  \tag{38}\\
& \mathcal{L}_{\mathrm{IPF}}(\theta) \text { is minimized } \Leftrightarrow D_{\mathrm{KL}}\left(q^{\phi} \| q^{\theta}\right)=0 \Leftrightarrow \forall t \in[0, T], Z_{t}^{\theta}+\widehat{Z}_{t}^{\phi}-\sigma \nabla \log q_{t}^{\phi}=0, \tag{39}
\end{align*}
$$

as implied by (27). If $\mathcal{L}_{\mathrm{TD}}(\phi)$ and $\mathcal{L}_{\mathrm{TD}}(\theta)$ are minimized, the following relations must also hold

$$
\begin{align*}
\mathrm{d} \widehat{Y}_{t}^{\phi} & =\left(\frac{1}{2}\left\|\widehat{Z}_{t}^{\phi}\right\|^{2}+\nabla \cdot\left(\sigma \widehat{Z}_{t}^{\phi}-f_{t}\right)+Z_{t}^{\theta^{\top}} \widehat{Z}_{t}^{\phi}-F_{t}\right) \mathrm{d} t+\widehat{Z}_{t}^{\phi \top} \mathrm{d} W_{t},  \tag{40}\\
\mathrm{~d} Y_{s}^{\theta} & =\left(\frac{1}{2}\left\|Z_{s}^{\theta}\right\|^{2}+\nabla \cdot\left(\sigma Z_{s}^{\theta}+f_{s}\right)+Z_{s}^{\theta^{\top}} \widehat{Z}_{s}^{\phi}-F_{s}\right) \mathrm{d} t+Z_{s}^{\theta \top} \mathrm{d} W_{s} . \tag{41}
\end{align*}
$$

Now, notice that the Fokker Plank equation of the parametrized forward SDE (25a) obeys

$$
\frac{\partial q_{t}^{\theta}}{\partial t}=-\nabla \cdot\left(q_{t}^{\theta}\left(f_{t}+\sigma Z_{t}^{\theta}\right)\right)+\frac{1}{2} \sigma^{2} \Delta q_{t}^{\theta}
$$

which implies that (c.f. Lemma 7),

$$
\begin{equation*}
\frac{\partial \log q_{t}^{\theta}}{\partial t}=-\nabla \cdot\left(f_{t}+\sigma Z_{t}^{\theta}\right)-\nabla \log q_{t}^{\theta^{\top}}\left(f_{t}+\sigma Z_{t}^{\theta}\right)+\frac{\sigma^{2}}{2}\left(\Delta \log q_{t}^{\theta}+\left\|\nabla \log q_{t}^{\theta}\right\|^{2}\right) \tag{42}
\end{equation*}
$$

Invoking Ito lemma yields:

$$
\begin{align*}
\mathrm{d} \log q_{t}^{\theta} & =\frac{\partial \log q_{t}^{\theta}}{\partial t} \mathrm{~d} t+\left[\nabla \log q_{t}^{\theta^{\top}}\left(f_{t}+\sigma Z_{t}^{\theta}\right)+\frac{\sigma^{2}}{2} \Delta \log q_{t}^{\theta}\right] \mathrm{d} t+\sigma \nabla \log q_{t}^{\theta^{\top}} \mathrm{d} W_{t} \\
& \stackrel{(42)}{=}\left[-\nabla \cdot\left(f_{t}+\sigma Z_{t}^{\theta}\right)+\sigma^{2} \Delta \log q_{t}^{\theta}+\frac{\sigma^{2}}{2}\left\|\nabla \log q_{t}^{\theta}\right\|^{2}\right] \mathrm{d} t+\sigma \nabla \log q_{t}^{\theta^{\top}} \mathrm{d} W_{t} \\
& =\left[-\nabla \cdot\left(f_{t}+\sigma Z_{t}^{\theta}-\sigma^{2} \nabla \log q_{t}^{\theta}\right)+\frac{1}{2}\left\|\sigma \nabla \log q_{t}^{\theta}\right\|^{2}\right] \mathrm{d} t+\sigma \nabla \log q_{t}^{\theta \top} \mathrm{d} W_{t} \\
& \stackrel{(*)}{=}\left[-\nabla \cdot\left(f_{t}+\sigma Z_{t}^{\theta}-\sigma\left(Z_{t}^{\theta}+\widehat{Z}_{t}^{\phi}\right)\right)+\frac{1}{2}\left\|Z_{t}^{\theta}+\widehat{Z}_{t}^{\phi}\right\|^{2}\right] \mathrm{d} t+\left(Z_{t}^{\theta}+\widehat{Z}_{t}^{\phi}\right)^{\top} \mathrm{d} W_{t} \\
& =\left[\nabla \cdot\left(\sigma \widehat{Z}_{t}^{\phi}-f_{t}\right)+\frac{1}{2}\left\|Z_{t}^{\theta}+\widehat{Z}_{t}^{\phi}\right\|^{2}\right] \mathrm{d} t+\left(Z_{t}^{\theta}+\widehat{Z}_{t}^{\phi}\right)^{\top} \mathrm{d} W_{t} \tag{43}
\end{align*}
$$

where $(*)$ is due to (38). Subtracting (40) from (43) yields

$$
\begin{equation*}
\mathrm{d} \log q_{t}^{\theta}-\mathrm{d} \widehat{Y}_{t}^{\phi}=\left(\frac{1}{2}\left\|Z_{t}^{\theta}\right\|^{2}+F_{t}\right) \mathrm{d} t+Z_{t}^{\theta^{\top}} \mathrm{d} W_{t} \tag{44}
\end{equation*}
$$

Now, using the fact that $Z_{\theta}:=\sigma \nabla Y_{\theta}$ and $\widehat{Z}_{\phi}:=\sigma \nabla \widehat{Y}_{\phi}$, we know that

$$
Z_{t}^{\theta}+\widehat{Z}_{t}^{\phi}-\sigma \nabla \log q_{t}^{\theta}=0 \Rightarrow Y_{t}^{\theta}+\widehat{Y}_{t}^{\phi}=\log q_{t}^{\theta}+c_{t}
$$

for some function $c_{t} \equiv c(t)$. Hence, (44) becomes

$$
\begin{equation*}
\mathrm{d} Y_{t}^{\theta}-\mathrm{d} c_{t}=\left(\frac{1}{2}\left\|Z_{t}^{\theta}\right\|^{2}+F_{t}\right) \mathrm{d} t+Z_{t}^{\theta^{\top}} \mathrm{d} W_{t} \tag{45}
\end{equation*}
$$

Now we prove that $\forall t \in(0, T), \mathrm{d} c_{t}=0$ by contradiction. First, notice that $c_{t}$ can be derived analytically as

$$
\begin{align*}
c_{t}= & Y_{t}^{\theta}+\widehat{Y}_{t}^{\phi}-\log q_{t}^{\theta} \\
= & \int_{0}^{t}\left(\mathrm{~d} Y_{\tau}^{\theta}+\mathrm{d} \widehat{Y}_{\tau}^{\phi}-\mathrm{d} \log q_{\tau}^{\theta}\right) \\
\stackrel{(*)}{=} & \int_{0}^{t}\left(\left(\frac{\partial Y_{\tau}^{\theta}}{\partial \tau}+\nabla Y_{\tau}^{\theta^{\top}}\left(f_{\tau}+\sigma Z_{\tau}^{\theta}\right)+\frac{\sigma^{2}}{2} \Delta Y_{\tau}^{\theta}\right)-\left(\frac{1}{2}\left\|Z_{\tau}^{\theta}\right\|^{2}+F_{\tau}\right)\right) \mathrm{d} \tau \\
& +\int_{0}^{t}\left(\sigma \nabla Y_{\tau}^{\theta^{\top}} \mathrm{d} W_{\tau}-Z_{\tau}^{\theta^{\top}} \mathrm{d} W_{\tau}\right) \\
\stackrel{(* *)}{=} & \int_{0}^{t}\left(\frac{\partial Y_{\tau}^{\theta}}{\partial t}-\left(-\nabla Y_{\tau}^{\theta^{\top}} f_{\tau}-\frac{1}{2}\left\|\sigma \nabla Y_{\tau}^{\theta}\right\|^{2}-\frac{\sigma^{2}}{2} \Delta Y_{\tau}^{\theta}+F_{\tau}\right)\right) \mathrm{d} \tau \tag{46}
\end{align*}
$$

where $\left({ }^{*}\right)$ invokes the following Ito lemma and substitutes (44),

$$
\mathrm{d} Y_{\tau}^{\theta}=\frac{\partial Y_{\tau}^{\theta}}{\partial \tau} \mathrm{d} \tau+\left[\nabla Y_{\tau}^{\theta^{\top}}\left(f_{\tau}+\sigma Z_{\tau}^{\theta}\right)+\frac{\sigma^{2}}{2} \Delta Y_{\tau}^{\theta}\right] \mathrm{d} t+\sigma \nabla Y_{\tau}^{\theta^{\top}} \mathrm{d} W_{\tau}
$$

and $\left({ }^{* *}\right)$ substitutes the definition $Z_{\tau}^{\theta}:=\sigma \nabla Y_{\tau}^{\theta}$. Equation (46) has an intriguing implication, as one can verify that its integrand is the residual of the parametrized $H J B Y_{\theta}=-u_{\theta} \approx-u$ (recall (7) and (8)). It is straightforward to see that, the residual shall also be preserved after the parametrized HJB is
expanded by Ito lemma w.r.t. the backward parametrized $\operatorname{SDE}(25 b)$. That is, the following equation similar to (45) must hold for the function $c_{s}:=c(T-t)$ :

$$
\mathrm{d} Y_{s}^{\theta}+\mathrm{d} c_{s}=\left(\frac{1}{2}\left\|Z_{s}^{\theta}\right\|^{2}+\nabla \cdot\left(f_{s}+\sigma Z_{s}^{\theta}\right)+Z_{s}^{\theta^{\top}} \widehat{Z}_{s}^{\phi}-F_{s}\right) \mathrm{d} t+Z_{s}^{\theta \top} \mathrm{d} W_{s}
$$

which contradicts (41). Hence, we must have $\mathrm{d} c_{s}=\mathrm{d} c_{t}=0$, and (45) becomes

$$
\begin{equation*}
\mathrm{d} Y_{t}^{\theta}=\left(\frac{1}{2}\left\|Z_{t}^{\theta}\right\|^{2}+F_{t}\right) \mathrm{d} t+Z_{t}^{\theta^{\top}} \mathrm{d} W_{t} \tag{47}
\end{equation*}
$$

In short, we have shown that, for the parametrized forward (25a) and backward (25b) SDEs, the fact that $(40,41)$ hold implies that $(47)$ holds, providing $\mathcal{L}_{\text {IPF }}$ is minimized. The exact same statement can be repeated to prove that

$$
\begin{equation*}
\mathrm{d} \widehat{Y}_{s}^{\phi}=\left(\frac{1}{2}\left\|\widehat{Z}_{s}^{\phi}\right\|^{2}+F_{s}\right) \mathrm{d} s+\widehat{Z}_{s}^{\phi \top} \mathrm{d} W_{s} \tag{48}
\end{equation*}
$$

Therefore, if the combined objectives are minimized, i.e., $(38,39,40,41)$ hold, the parametrized functions $\left(Y_{\theta}, Z_{\theta}, \widehat{Y}_{\phi}, \widehat{Z}_{\phi}\right)$ satisfy $(25,40,47,41,48)$, i.e., they satisfy the $\operatorname{FBSDE}$ systems $(11,12)$ in Theorem 2.

## A. 4 Additional Derivations \& Remarks in Sec. 3 and 4

## A.4.1 Hopf-Cole transform

Recall the Hopf-Cole transform

$$
\Psi(x, t):=\exp (-u(x, t)), \quad \widehat{\Psi}(x, t):=\rho(x, t) \exp (u(x, t))
$$

Standard ordinary calculus yields

$$
\begin{array}{ll}
\nabla \Psi=-\exp (-u) \nabla u, & \Delta \Psi=\exp (-u)\left[\|\nabla u\|^{2}-\Delta u\right] \\
\nabla \widehat{\Psi}=\exp (u)(\rho \nabla u+\nabla \rho), & \Delta \widehat{\Psi}=\exp (u)\left[\rho\|\nabla u\|^{2}+2 \nabla \rho^{\top} \nabla u+\Delta \rho+\rho \Delta u\right]
\end{array}
$$

Hence, we have

$$
\begin{align*}
& \frac{\partial \Psi}{\partial t}=\exp (-u)\left(-\frac{\partial u}{\partial t}\right) \\
& \stackrel{(7)}{=} \exp (-u)\left(-\frac{1}{2}\|\sigma \nabla u\|^{2}+\nabla u^{\top} f+\frac{1}{2} \sigma^{2} \Delta u+F\right) \\
& \stackrel{(49)}{=}-\frac{1}{2} \sigma^{2} \Delta \Psi-\nabla \Psi^{\top} f+F \Psi  \tag{51}\\
& \frac{\partial \widehat{\Psi}}{\partial t}=\exp (u)\left(\frac{\partial \rho}{\partial t}+\rho \frac{\partial u}{\partial t}\right) \\
& \stackrel{(7)}{=} \exp (u)\left(\left(\nabla \cdot\left(\rho\left(\sigma^{2} \nabla u-f\right)\right)+\frac{1}{2} \sigma^{2} \Delta \rho\right)+\rho\left(\frac{1}{2}\|\sigma \nabla u\|^{2}-\nabla u^{\top} f-\frac{1}{2} \sigma^{2} \Delta u-F\right)\right) \\
&=\exp (u)\left(\sigma^{2}\left(\rho \Delta u+\nabla \rho^{\top} \nabla u+\frac{1}{2} \Delta \rho+\frac{\rho}{2}\|\nabla u\|^{2}-\frac{\rho}{2} \Delta u\right)-\nabla \rho^{\top} f-\rho \nabla \cdot f-\rho \nabla u^{\top} f-\rho F\right) \\
& \stackrel{(50)}{=} \frac{1}{2} \sigma^{2} \Delta \widehat{\Psi}-\nabla \widehat{\Psi}^{\top} f-\widehat{\Psi} \nabla \cdot f-\widehat{\Psi} F, \tag{52}
\end{align*}
$$

which yields (9) by noticing that $\nabla \cdot(\widehat{\Psi} f)=\nabla \widehat{\Psi}^{\top} f+\widehat{\Psi} \nabla \cdot f$.

## A.4.2 Remarks on convergence

The alternating optimization scheme proposed in Alg. (1) can be compactly presented as $\min _{\phi} D_{\mathrm{KL}}\left(q^{\theta} \| q^{\phi}\right)+\mathbb{E}_{q^{\theta}}\left[\mathcal{L}_{\mathrm{TD}}(\phi)\right]$ and $\min _{\theta} D_{\mathrm{KL}}\left(q^{\phi} \mid q^{\theta}\right)+\mathbb{E}_{q^{\phi}}\left[\mathcal{L}_{\mathrm{TD}}(\theta)\right]$. Despite that the procedure seems to resemble IPF, which optimizes between $\min _{\phi} D_{\mathrm{KL}}\left(q^{\phi} \mid q^{\theta}\right)$ and $\min _{\theta} D_{\mathrm{KL}}\left(q^{\theta} \mid q^{\phi}\right)$, we stress that they differ from each other in that the the KLs are constructed with different directions.

In cases where the TD objectives are discarded, prior work [24] has proven that minimizing the forward KLs admit similar convergence to standard IPF (which minimizes the reversed KLs). This is essentially the key to developing scalable methods, since the parameter being optimized (e.g., $\theta$ in $D_{\mathrm{KL}}\left(q^{\phi} \mid q^{\theta}\right)$ ) in forward KLs differs from the parameter used to sample expectation (e.g., $\mathbb{E}_{q^{\phi}}$ ). Therefore, the computational graph of the SDEs can be dropped, yielding a computationally much efficient framework. These advantages have been adopted in [24,27] and also this work for solving higher-dimensional problems.

However, when we need TD objectives to enforce the MF structure, as appeared in all the MFGs in this work, the combined objective does not correspond to IPF straightforwardly. Despite that the alternating procedure in Alg. (1) is mainly inspired by prior SB methods [24, 27], the training process of DeepGSB is perhaps closer to TRPO [55], which iteratively updates the policy using the off-policy samples generated from the previous stage: $\pi^{(i+1)}=\arg \min _{\pi} D_{\mathrm{KL}}\left(\pi^{(i)} \| \pi\right)+\mathbb{E}_{\pi^{(i)}}[\mathcal{L}(\pi)]$. TRPO is proven to enjoy monotonic improvement over iterations (i.e., local convergence).

## A.4.3 Functional derivative of MF potential functions

Given a functional $\mathcal{F}: \mathcal{P}\left(\mathbb{R}^{d}\right) \rightarrow \mathbb{R}$ on the space of probability measures, its functional derivative $F(x, \rho):=\frac{\delta \mathcal{F}(\rho)}{\delta \rho}(x)$ satisfies the following equation

$$
\lim _{h \rightarrow 0} \frac{\mathcal{F}(\rho+h w)-\mathcal{F}(\rho)}{h}=\int_{\mathbb{R}^{d}} F(x, \rho) w(x) \mathrm{d} x
$$

for any function $w \in L^{2}\left(\mathbb{R}^{d}\right)$. Hence, the derivative of the entropy MF functional $\mathcal{F}_{\text {entropy }}:=$ $\int_{\mathbb{R}^{d}} \rho(x) \log \rho(x) \mathrm{d} x$ can be derived as

$$
\begin{align*}
& \lim _{h \rightarrow 0} \frac{1}{h}\left(\mathcal{F}_{\text {entropy }}(\rho+h w)-\mathcal{F}_{\text {entropy }}(\rho)\right) \\
= & \lim _{h \rightarrow 0} \frac{1}{h}\left(\int_{\mathbb{R}^{d}}\left(h w(x) \log \rho(x)+\rho(x) \frac{h w(x)}{\rho(x)}+\mathcal{O}\left(h^{2}\right)\right) \mathrm{d} x\right) \\
= & \int_{\mathbb{R}^{d}}(w(x) \log \rho(x)+w(x)) \mathrm{d} x=\int_{\mathbb{R}^{d}} \underbrace{(\log \rho(x)+1)}_{:=F_{\text {entropy }}(x, \rho)} w(x) \mathrm{d} x . \tag{53}
\end{align*}
$$

Similarly, consider the congestion MF functional $\mathcal{F}_{\text {congestion }}:=\int_{\mathbb{R}^{d}} \int_{\mathbb{R}^{d}} \frac{1}{\|x-y\|^{2}+1} \rho(x) \rho(y) \mathrm{d} x \mathrm{~d} y$. Its derivation can be computed by

$$
\begin{align*}
& \lim _{h \rightarrow 0} \frac{1}{h}\left(\mathcal{F}_{\text {congestion }}(\rho+h w)-\mathcal{F}_{\text {congestion }}(\rho)\right) \\
= & \lim _{h \rightarrow 0} \frac{1}{h}\left(\int_{\mathbb{R}^{d}} \int_{\mathbb{R}^{d}} \frac{1}{\|x-y\|^{2}+1}\left(\rho(x) h w(y)+h w(x) \rho(y)+\mathcal{O}\left(h^{2}\right)\right) \mathrm{d} x \mathrm{~d} y\right) \\
= & \int_{\mathbb{R}^{d}} \int_{\mathbb{R}^{d}} \frac{1}{\|x-y\|^{2}+1}(\rho(x) w(y)+w(x) \rho(y)) \mathrm{d} x \mathrm{~d} y \\
= & \int_{\mathbb{R}^{d}} \underbrace{\int_{\mathbb{R}^{d}} \frac{2}{\|x-y\|^{2}+1} \mathrm{~d} y}_{:=F_{\text {congestion }}(x, \rho)} w(x) \mathrm{d} x . \tag{54}
\end{align*}
$$

We hence conclude the expressions of $F_{\text {entropy }}$ and $F_{\text {congestion }}$ in (17).

## A. 5 Experiment Details

## A.5.1 Setup

Hyperparameters Table 6 summarizes the hyperparameters in each MFG, including the dimension $d$ of the state space, the diffusion scalar $\sigma$, the time horizon $T$, the discretized time step $\delta t$ (and $\delta s$ ), the MF base drift $f(x, \rho)$, the MF interaction $F(x, \rho)$, and the mean/covariance of the boundary distributions $\rho_{0}$ and $\rho_{\text {target }}$ ( note that all MFGs adopt Gaussians as their boundary distributions ). Note that in the 1000 -dimensional opinion MFG, we multiply the polarized dynamic $\bar{f}_{\text {polarize }}$ by 6 to

Table 6: Hyperparameters in each MFG. Note that $\mathbf{0} \in \mathbb{R}^{d}$ denotes zero vector, $\boldsymbol{I} \in \mathbb{R}^{d \times d}$ denotes identity matrix, and $\operatorname{diag}(\boldsymbol{v}) \in \mathbb{R}^{d \times d}$, where $\boldsymbol{v} \in \mathbb{R}^{d}$, denotes diagonal matrix.

|  | GMM | V-neck | S-tunnel | Opinion |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $d$ | 2 | 2 | 2 | 2 | 1000 |
| $\sigma$ | 1 | 1 | 1 | 0.1 | 0.5 |
| $T$ | 1 | 2 | 3 | 3 | 3 |
| $\delta t$ | 0.01 | 0.01 | 0.01 | 0.01 | 0.006 |
| $f(x, \rho)$ | $[0,0]^{\top}$ | $[6,0]^{\top}$ | $[6,0]^{\top}$ | $\bar{f}_{\text {polarize }}$ | $6 \cdot \bar{f}_{\text {polarize }}$ |
| Diffusion steps | 100 | 200 | 300 | 300 | 500 |
| $K^{7}$ | 250 | 250 | 500 | 100 | 250 |
| Alternating stages ${ }^{8}$ | 40 | 40 | 30 | 40 | 90 |
| Total training steps | 20k | 20k | 30k | 8 k | 45k |
| Mean of $\rho_{0}$ | 0 | $\left[\begin{array}{c}-7 \\ 0\end{array}\right]$ | $\left[\begin{array}{c}-11 \\ -1\end{array}\right]$ | 0 | 0 |
| Mean of $\rho_{\text {target }}$ | $\begin{gathered} e^{16 \cdot\left(\frac{\pi}{4}\right) i}, \\ i \in\{0, \cdots, 7\} \end{gathered}$ | $\left[\begin{array}{l}7 \\ 0\end{array}\right]$ | $\left[\begin{array}{c}11 \\ 1\end{array}\right]$ | 0 | 0 |
| Covariance of $\rho_{0}$ | $\boldsymbol{I}$ | 0.2 I | 0.51 | $\operatorname{diag}\left(\left[\begin{array}{c}0.5 \\ 0.25\end{array}\right]\right)$ | $\operatorname{diag}\left(\left[\begin{array}{c} 4 \\ 0.25 \\ \vdots \\ 0.25 \end{array}\right]\right)$ |
| Covariance of $\rho_{\text {target }}$ | $I$ | $0.2 \boldsymbol{I}$ | $0.5 I$ | 3 I | 3 I |

ensure that the high-dimensional dynamics yield polarization within the time horizon. Meanwhile, a smaller step size $\delta t=0.006$ is adopted so that the discretization error from the relatively large drift is mitigated. As mentioned in Sec. 4, we adopt zero and constant base drift $f$ respectively for GMM and V-neck/S-tunnel.

Training All experiments are conducted on 3 TITAN RTXs and 1 TITAN V100, where the V100 is located on the Amazon Web Service (AWS). We use the multi-step TD targets in (15) for all experiments and adopt huber norm for the TD loss in (16). As for the FK consistency loss $\mathcal{L}_{\mathrm{FK}}$, we use $\ell 1$ norm for GMM and opinion MFGs, and huber norm for the rest.

Network architecture All networks $\left(Y_{\theta}, Z_{\theta}, \widehat{Y}_{\phi}, \widehat{Z}_{\phi}\right)$ take $(x, t)$ as inputs and follow

$$
\text { out }=\text { out_mod }\left(\mathrm{x}_{-} \bmod (x)+\mathrm{t}_{-} \bmod (\text { timestep_embedding }(t))\right), ~_{\text {, }}
$$

where timestep_embedding $(\cdot)$ is the standard sinusoidal embedding.
For crowd navigation MFGs, these modules consist of 2 to 4 fully-connected layers (Linear) followed by the Sigmoid Linear Unit (SiLU) activation functions [70], i.e.,

$$
\begin{aligned}
\text { t_mod } & =\text { Linear } \rightarrow \text { SiLU } \rightarrow \text { Linear } \\
\text { x_mod } & =\text { Linear } \rightarrow \text { SiLU } \rightarrow \text { Linear } \rightarrow \text { SiLU } \rightarrow \text { Linear } \rightarrow \text { SiLU } \rightarrow \text { Linear } \\
\text { out_mod } & =\text { Linear } \rightarrow \text { SiLU } \rightarrow \text { Linear } \rightarrow \text { SiLU } \rightarrow \text { Linear }
\end{aligned}
$$

As for 1000-dimensional opinion MFG, we keep the same t_mod and out_mod but adopt residual networks with 5 residual blocks for x_mod. For DeepGSB-ac, we set the hidden dimension of Linear to 256 and 128 respectively for the policy networks $\left(Z_{\theta}, \widehat{Z}_{\phi}\right)$ and the critic networks $\left(Y_{\theta}, \widehat{Y}_{\phi}\right)$, whereas for DeepGSB-c, we set the hidden dimension of Linear to 200 for the critic networks $\left(Y_{\theta}, \widehat{Y}_{\phi}\right)$.

[^0]Implementation of prior methods [14-16] All of our experiments are implemented with PyTorch [71]. Hence, we re-implement the method in Ruthotto et al. [14] by migrating their Julia codebase ${ }^{9}$ to PyTorch. As for Lin et al. [15], their official PyTorch implementation is publicly available. ${ }^{10}$ Finally, we implement Chen [16] by ourselves. Since prior methods [14-16] were developed for a smaller class of MFGs compared to our DeepGSB (recall Table 1), we need to relax the setup of the MFG in order for them to yield reasonable results in Fig. 5 and 7. Specifically, we soften the obstacle costs, so that $[14,15]$ can differentiate them properly, and keep the same KL penalty at $u(x, T) \approx D_{\mathrm{KL}}\left(\rho(x, T) \| \rho_{\text {target }}(x)\right)$ as adopted in $[14,15]$. We stress that neither of the methods [14, 15] works well with the discontinuous $F_{\text {obstacle }}$ in (17). Finally, we discretize the 2-dimensional state space of GMM into a $40 \times 40$ grid with 50 time steps for [16]. We note that the complexity of [16] scales as $\mathcal{O}\left(\tilde{T} D^{2}\right)$, where $\tilde{T}$ and $D$ are respectively the number of time and spatial grids, i.e., $\tilde{T}=50$ and $D=1600$.

Evaluation We approximate the Wasserstein distance with the Sinkhorn divergence using the geomloss package. ${ }^{11}$ The Sinkhorn divergence interpolates between Wasserstein (blur $=0$ ) and kernel (blur $=\infty$ ) distance given the hyperparameter blur. We set blur $=0.05$ in Table 3.

## A.5.2 Additional experiments



Figure 7: Same setup as in Fig. 5 except for DeepGSB-c. This figure is best viewed in color.

Figures 7 and 8 reports the results for DeepGSB-c. On crowd navigation MFGs, the population snapshots guided by DeepGSB-c are visually indistinguishable from DeepGSB-ac (see Fig. 7 vs. 5) despite the visual difference in their contours. As for 1000-dimensional opinion MFG, both DeepGSBc and DeepGSB-ac are able to guild the population opinions toward desired $\rho_{\text {target }}$ without the entropy interaction $F_{\text {entropy }}$. Figure 8 reports the results of DeepGSB-c in such cases. We note, however, that when $F_{\text {entropy }}$ is enabled, DeepGSB-ac typically performs better than DeepGSB-c in terms of convergence to $\rho_{\text {target }}$ and training stability.

[^1]

Figure 8: (a) Visualization of polarized dynamics $\bar{f}_{\text {polarize }}$ in 2- and 1000-dimensional opinion space, where the directional similarity [3] counts the histogram of cosine angle between pairwise opinions at the terminal distribution $\rho_{T}$. (b) DeepGSB-c guides $\rho_{T}$ to approach moderated distributions, hence depolarizes the opinion dynamics. Note that we adopt $F:=0$ for DeepGSB-c. We use the first two principal components to visualize $d=1000$.


Figure 9: DeepGSB-ac trained without access to the initial and target distributions, i.e., without $\mathrm{TD}_{0}$ and $\widehat{\mathrm{TD}}_{0}$. In this case, we compute $\mathcal{L}_{\mathrm{TD}}$ with the single-step formulation in (14).


[^0]:    ${ }^{7}$ We note that, unlike SB-FBSDE [27], the number of training iterations at each stage (i.e., the $K$ in Alg. 1) is kept fixed throughout training.
    ${ }^{8}$ Here, we refer one alternating stage to a complete cycling through $2 K$ training iterations in Alg. 1.

[^1]:    ${ }^{9}$ https://github.com/EmoryMLIP/MFGnet.jl. The repository is licensed under MIT License.
    ${ }^{10}$ https://github.com/atlin23/apac-net. The repository does not specify licenses.
    ${ }^{11}$ https://github.com/jeanfeydy/geomloss. The repository is licensed under MIT License.

