Target Mirror Descent: A Unifying Framework for Solving Monotone Variational Inequalities

Yu-Wen Chen, Can Kizilkale, Murat Arcak

***Abstract*—It is well known that mirror descent may diverge or cycle on merely monotone variational inequalities. In this paper, we propose *Target Mirror Descent* (TMD), a unified framework that stabilizes monotone flows via a target point correction mechanism in the dual update. By appropriate design choices, the TMD framework recovers the proximal point algorithm, extragradient methods, splitting methods, Brown-von Neumann-Nash dynamics, forward-backward-forward dynamics, and discounted mirror descent as special instances. Thus, we establish a unified perspective on these landmark algorithms and their convergence. Beyond unification, we leverage the TMD framework to correct an equilibrium misalignment in discounted mirror descent and to generalize its higher-order extension beyond interior solutions. Moreover, a key structural feature of TMD is the explicit decoupling of the mirror map from the target determination, which enables *geometric ensembles*: multiple TMD instances solve the same problem in parallel using distinct mirror maps, while sharing a common dual update. We show that such an ensemble rigorously reduces to a single TMD with a synthesized mirror map, and thus inherits these convergence guarantees.**

***Index Terms*—Optimization algorithms, variational inequalities, stability of nonlinear systems, game theory**

I. Introduction

VARIATIONAL inequalities (VIs) provide a powerful mathematical framework for modeling equilibrium problems in constrained systems, encompassing optimization, game theory, and physical systems as special cases. Unlike scalar minimization, VIs are governed by a vector field F that may lack a potential function, arising naturally in multi-agent games, traffic equilibria, and resource allocation.

A central challenge is designing algorithms that converge reliably when F is merely monotone, i.e., monotone but not necessarily strictly monotone. It is known that mirror descent may cycle or diverge in this regime [1]. Reference [2] shows the ergodic (time-averaged) convergence of mirror descent, while last-iterate convergence requires more sophisticated techniques. To this end, several stabilizing algorithms have been developed, each with its own convergence analysis: the proximal point algorithm [3], extragradient methods [1], [4], [5], splitting methods [6]–[8], Brown-von Neumann-Nash (BNN) dynamics [9], forward-backward-forward dynamics [10], [11], and (improved) discounted mirror descent [12]. While [13] unifies extragradient and optimistic gradient methods as approximations of the proximal point method, a broader unifying perspective is missing. In particular, the connections between game-theoretic dynamics such as BNN and classical optimization algorithms have not been fully explored.

Yu-Wen Chen, Can Kizilkale, and Murat Arcak are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA, USA. {yuwen_chen, cankizilkale, arcak}@berkeley.edu.

All of these algorithms operate with a single fixed geometric structure, leaving open the question of whether multiple geometries can be combined. In machine learning, ensemble methods combine multiple models to improve robustness. The underlying intuition, often called the *wisdom of crowds*, is that a diverse collection of models captures complementary specialization that no single model can provide alone. In the context of mirror descent, this motivates a *geometric ensemble*: multiple algorithms, each equipped with a geometrically distinct mirror map, solving the same VI problem in parallel.

The main contributions of this paper are as follows:

- We propose *Target Mirror Descent* (TMD), a unified framework that stabilizes monotone flows via a target point correction mechanism in the dual update, with a rigorous convergence analysis.
- We show that, by appropriate design choices, TMD subsumes proximal point algorithms, extragradient methods, splitting methods, BNN dynamics, forward-backward-forward dynamics, discounted mirror descent and its higher-order extension as special instances, establishing a unified structure and analysis across all of them.
- A key structural feature of TMD is the explicit decoupling of the mirror map from the target determination. This separation enables *geometric ensembles*: multiple TMD instances with distinct mirror maps, whose dual updates are driven by a shared ensemble state. We show that the ensemble rigorously reduces to a single TMD with a synthesized mirror map, inheriting convergence guarantees.

The remainder of this paper is as follows. Section II introduces variational inequalities and mirror descent. Section III presents the TMD framework, its convergence analysis, and the unifying perspective. Section IV develops the geometric ensemble. Finally, Section V presents numerical validation.

Notation: $\langle \cdot, \cdot \rangle$ denotes the inner product, $\top$ the transpose, ∇f the gradient of f, $\circ$ the composition of functions, $\mathrm{Im}(f)$ the image of f, $\mathrm{int}(\mathcal{X})$ the interior of $\mathcal{X}$, $\mathbb{R}^n_{>0}$ the positive orthant, $\|\cdot\|$ a norm, $\Delta^n = \{x \in \mathbb{R}^n : x \succeq 0, \mathbf{1}^\top x = 1\}$ the n-dimensional standard simplex, and $\mathcal{N}_{\mathcal{X}}(x) = \{v \in \mathbb{R}^n : \langle v, y - x \rangle \leq 0, \forall y \in \mathcal{X} \subseteq \mathbb{R}^n\}$ the normal cone of $\mathcal{X}$ at x.

II. Preliminaries

A. Variational Inequalities

Definition 1 (Variational Inequality): Given a set $\mathcal{X} \subseteq \mathbb{R}^n$ and a map $F : \mathcal{X} \to \mathbb{R}^n$, an $x^* \in \mathcal{X}$ solves the variational inequality, denoted by $\mathrm{VI}(\mathcal{X}, F)$, if $F(x^*)^\top(x-x^*) \geq 0, \forall x \in \mathcal{X}$. The set of all solutions is denoted as $\mathrm{SOL}(\mathcal{X}, F)$.

Some common properties studied for F are as follows.

Definition 2 (Monotonicity): A map $F : \mathcal{X} \to \mathbb{R}^n$ is *pseudo-monotone* if $F(y)^\top(x-y) \geq 0 \Rightarrow F(x)^\top(x-y) \geq 0$; *monotone* if $(F(x) - F(y))^\top(x-y) \geq 0$; and *$\sigma$-strongly monotone* if $(F(x) - F(y))^\top(x-y) \geq \sigma\|x-y\|^2$, for some $\sigma > 0$ and norm $\|\cdot\|$, for all $x, y \in \mathcal{X}$.

Definition 3 (Variational Stability, [14, Definition 2.3]): A point $\bar{x} \in \mathcal{X}$ is variationally stable[1] (VS) w.r.t. $F : \mathcal{X} \to \mathbb{R}^n$, if $F(x)^\top(x-\bar{x}) \geq 0, \forall x \in \mathcal{X}$. A set $\mathcal{S} \subseteq \mathcal{X}$ is VS w.r.t. F, if the inequality holds $\forall \bar{x} \in \mathcal{S}$.

Note that strong monotonicity implies monotonicity, which in turn implies pseudo-monotonicity. Moreover, if F is pseudo-monotone, then $\mathrm{SOL}(\mathcal{X}, F)$ is VS w.r.t. F.

B. Mirror Descent

Mirror descent solves strongly monotone variational inequalities. It generalizes gradient descent by operating in a dual space induced by a strongly convex or Legendre function $h : \mathcal{X} \to \mathbb{R}$, where $\mathcal{X} \subseteq \mathbb{R}^n$ is nonempty, closed, and convex. Denoting the convex conjugate of h by h^*, the map $\nabla h : \mathcal{X} \to \mathbb{R}^n$ sends primal iterates to the dual space, where gradient updates are performed; $\nabla h^* : \mathbb{R}^n \to \mathcal{X}$ then maps back to the primal space. In this paper, we refer both ∇h and ∇h^* as *mirror maps*, which have an important relation[2]: $(\nabla h)^{-1} = \nabla h^*$. The continuous version of mirror descent is

$$\dot{z}(t) = -\eta F(x(t)), \quad x(t) = \nabla h^*(z(t)), \tag{1}$$

where x and z are the primal and dual states, respectively.

III. Target Mirror Descent: A Unified Framework

In this paper, we consider variational inequalities with closed and convex $\mathcal{X} \subseteq \mathbb{R}^n$ and continuous F. As noted in [1] and [6, Chapter 12.1], gradient (or mirror) descent may diverge or cycle when applied to merely monotone F, which motivates our proposed TMD framework.

In Section III-A, we introduce the *Target Mirror Descent* (TMD) framework and provide the convergence analysis in Section III-B. Beyond stability, we demonstrate in Section III-C that TMD, as a unifying framework, subsumes many landmark algorithms. Throughout, we assume well-posedness:

Assumption 1: $\mathrm{SOL}(\mathcal{X}, F) \neq \emptyset$.

A. Target Mirror Descent (TMD) dynamics

Our main idea is to provide feedback correction to mirror descent (1). Specifically, for a given state x, we compute a *target point* $T(x)$. Then, a design function S maps both x and $T(x)$ to the dual space, where the dual discrepancy $S(T(x)) - S(x)$ provides a corrected update direction. Mathematically, the proposed *target mirror descent (TMD)* dynamics is

$$\dot{z}(t) = \alpha\,(S(T(x(t))) - S(x(t))) - \beta\Phi(x(t)) \tag{2a}$$
$$x(t) = \nabla h^*(z(t)), \tag{2b}$$

where $\alpha > 0$ and $\beta \geq 0$ are parameters, $\nabla h^* : \mathbb{R}^n \to \mathcal{X}$ is the mirror map introduced in Section II-B, and $T : \mathcal{X} \to \mathcal{X}$, $S : \mathcal{X} \to \mathbb{R}^n$, and $\Phi : \mathcal{X} \to \mathbb{R}^n$ are design functions satisfying:

- [C1] S is σ-strongly monotone w.r.t. some norm $\|\cdot\|$.
- [C2] There exists an $\bar{x} \in \mathcal{X}$ which is VS w.r.t. Φ.
- [C3] $\mathrm{SOL}(\mathcal{X}, \Phi) \subseteq \mathrm{SOL}(\mathcal{X}, F)$.
- [C4] $T = (S + \Phi + \mathcal{N}_\mathcal{X})^{-1} \circ S$ is single-valued.

Here, $(S + \Phi + \mathcal{N}_\mathcal{X})(x) = S(x) + \Phi(x) + \mathcal{N}_\mathcal{X}(x)$ refers to the Minkowski sum. By embedding $\mathcal{N}_\mathcal{X}$ directly into the formulation of T, the system implicitly enforces feasibility, acting as a projection that ensures $\mathrm{Im}(T) \subseteq \mathcal{X}$.

The selection of S, Φ, and T offers significant flexibility, but the conditions C1–C4 must be satisfied. C1 requires the design function S to provide a strongly monotone structure for the dual update and guarantee the target mapping T to be well-defined. C2 stipulates a variationally stable equilibrium for the design function Φ, ensuring long-term convergence. C3 guarantees that every solution to the surrogate Φ is simultaneously a valid solution to the target problem F. Finally, C4 ensures that a valid target point can always be calculated. Moreover, each point x satisfying $x = T(x)$ must be a solution to $\mathrm{SOL}(\mathcal{X}, \Phi)$. In the following, we provide a sufficient condition for C4; weaker conditions can be found in [6, Chapter 2].

Lemma 1 ([6, Theorem 2.3.3]): If $S+\Phi$ is strongly monotone, then T is single-valued.

Example 1 (L-Lipschitz continuous & pseudo-monotone F): Pick a σ-strongly monotone S and $\Phi = \eta F$ with $0 < \eta < \frac{\sigma}{L}$, then $S + \Phi$ is strongly monotone, rendering T well-defined by Lemma 1. Thus, all C1–C4 are satisfied.

The formulation admits a graphical interpretation, shown in Figure 1. We emphasize that TMD is a *framework*: concrete algorithms arise by instantiating the design functions S, T, and Φ with specific choices, as demonstrated in Section III-C.

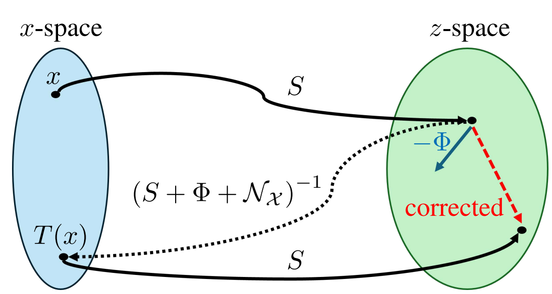

Fig. 1. Graphical illustration of TMD (2). The target point is $\boldsymbol{T(x)}$ and the corrected direction $\boldsymbol{S(T(x)) - S(x)}$ is marked by the red dashed line.

A key feature of TMD dynamics (2) is the decoupling of the mirror map ∇h^* from the design function S in the dual update (2a). Essentially, ∇h^*, as in mirror descent, encodes the problem geometry, while the design function S is selected to generate corrective feedback in the dual space. As we show

[1] It is also called evolutionary stability [9, Chapter 3.3.4].

[2] Refer to [12, Lemma 2 & Lemma 3] for detailed descriptions.

in Section IV, this decoupling is essential to the *geometric ensemble* technique: multiple TMD instances maintain distinct mirror maps while sharing a common dual update.

B. Convergence analysis for TMD

In this subsection, we establish the convergence of TMD.

Theorem 1: Under conditions C1–C4, the trajectories generated by TMD (2) asymptotically approach $\mathrm{SOL}(\mathcal{X}, \Phi)$.

Proof: Denote $D_h(s,t) = h(s) - h(t) - \langle \nabla h(t), s-t \rangle$ as the Bregman divergence, which is positive whenever $s \neq t$. Consider the Lyapunov function candidate $V(x) = D_h(\bar{x}, x)$:

$$
\begin{aligned}
\frac{d}{dt}V(x(t)) &= -\left\langle \nabla^2 h(x(t))\dot{x}(t), \bar{x} - x(t) \right\rangle && \text{(3a)}\\
&= -\left\langle \dot{z}(t), \bar{x} - x(t) \right\rangle && \text{(3b)}\\
&\leq -\alpha \left\langle (S \circ T - S)(x(t)), \bar{x} - T(x(t)) \right\rangle \\
&\quad - \alpha \left\langle (S \circ T - S)(x(t)), T(x(t)) - x(t) \right\rangle && \text{(3c)}\\
&= -\alpha \left\langle \Phi(T(x(t))), T(x(t)) - \bar{x} \right\rangle + \alpha \langle \theta(t), \bar{x} - T(x(t)) \rangle \\
&\quad - \alpha \left\langle S(T(x(t))) - S(x(t)), T(x(t)) - x(t) \right\rangle && \text{(3d)}\\
&\leq -\alpha\sigma \|T(x(t)) - x(t)\|^2, && \text{(3e)}
\end{aligned}
$$

where $\theta(t) \in \mathcal{N}_{\mathcal{X}}(T(x(t)))$. Note that (3b) follows from $(\nabla h)^{-1} = \nabla h^*$ from Section II-B and differentiating $z = \nabla h(x)$, (3c) from C2, (3d) from C4, and (3e) from the property of the normal cone, C1, and C2. By Barbalat's lemma, $x(t)$ asymptotically approaches the set of fixed points of T.

It remains to show that the fixed points of T are contained in $\mathrm{SOL}(\mathcal{X}, \Phi)$. From C4, with $y = Tx$, we have $S(x) - S(y) - \Phi(y) \in \mathcal{N}_{\mathcal{X}}(y)$. By the definition of the normal cone, we get $\langle S(x) - S(y) - \Phi(y), u - y \rangle \leq 0, \forall u \in \mathcal{X}$. For any fixed point x^* of T, substituting $x = x^*$ and $y = Tx^* = x^*$ into it yields $\langle \Phi(x^*), u - x^* \rangle \geq 0, \forall u \in \mathcal{X}$. Thus, the theorem follows. ∎

Asymptotic approach to a solution set does not guarantee convergence to a specific point within it. To bridge this gap, we consider the stronger condition given as follows.

[C2$^+$] $\mathrm{SOL}(\mathcal{X}, \Phi)$ is VS w.r.t. Φ.

One example satisfying C2$^+$ is the pseudo-monotone F; see the discussion after Definition 3.

Corollary 1: Under conditions C1, C2$^+$, C3, and C4, the trajectories generated by TMD (2) converge to $x^* \in \mathrm{SOL}(\mathcal{X}, \Phi)$.

Proof: By C2$^+$, the Lyapunov function $V(x) = D_h(\bar{x}, x)$ is valid for *any* $\bar{x} \in \mathrm{SOL}(\mathcal{X}, \Phi)$. In other words, $x(t)$ is Bregman-Fejér monotone w.r.t. $\mathrm{SOL}(\mathcal{X}, \Phi)$ [15]. Thus, by [15, Theorem 4.11], $x(t)$ converges to a point in the set. ∎

The term $\|T(x) - x\|^2$ in (3e) is positive whenever $x \notin \mathrm{SOL}(\mathcal{X}, \Phi)$. It can compensate for a mild violation of C2, e.g., when no $\bar{x}$ is perfectly VS w.r.t. Φ, leading to the following relaxed condition C2$^-$. Note that C2 implies C2$^-$ trivially.

[C2$^-$] There exists an $\bar{x} \in \mathcal{X}$ s.t. $\forall x \notin \mathrm{SOL}(\mathcal{X}, \Phi)$,

$$
\begin{aligned}
\alpha \left(\sigma \|T(x) - x\|^2 + \langle \Phi(T(x)), T(x) - \bar{x} \rangle \right) \\
+ \beta \langle \Phi(x), x - \bar{x} \rangle > 0. \quad \text{(4)}
\end{aligned}
$$

Corollary 2: Under conditions C1, C2$^-$, C3, and C4, the trajectories generated by TMD (2) asymptotically approach $\mathrm{SOL}(\mathcal{X}, \Phi)$. If C2$^-$ holds for all $\bar{x} \in \mathrm{SOL}(\mathcal{X}, \Phi)$, then the trajectories converge to $x^* \in \mathrm{SOL}(\mathcal{X}, \Phi)$.

C. TMD as a unifying framework

A central contribution of TMD is its ability to subsume several landmark algorithms as special instances. To facilitate connections with these iterative algorithms, we provide a discrete-time representation of (2a) in terms of x:[3]

$$
x_{t+1} = \nabla h^* (\nabla h(x_t) + \alpha (S \circ T - S)(x_t) - \beta \Phi(x_t)) . \quad (5)
$$

Unless otherwise noted, we assume F Lipschitz continuous and pseudo-monotone and ∇h strongly monotone. Let $\eta > 0$. In each following part, we first state the original algorithm and then provide a proposition showing that TMD recovers it.

1) (Nonlinear) Proximal Point Algorithms (PPA) [3]:

$$
x_{t+1} = (\nabla h + \eta F + \mathcal{N}_{\mathcal{X}})^{-1} (\nabla h(x_t)) . \quad (6)
$$

Proposition 1: Let $\alpha = 1$, $\beta = 0$, $S = \nabla h$, and $\Phi = \eta F$. Then, (5) recovers (6), and convergence follows from Corollary 1.

Proof: There exists a sufficiently small η s.t. $\nabla h + \eta F$ is strongly monotone. Thus, $T = (\nabla h + \eta F + \mathcal{N}_{\mathcal{X}})^{-1} \circ \nabla h$ is well-defined. All C1–C4 and C2$^+$ are satisfied. Substituting the selections into (5) yields (6). ∎

2) Extragradient methods (EG) or Mirror-Prox [1], [4]:

$$
\begin{aligned}
\omega_t &= \arg\min_{y \in \mathcal{X}} \; \{ \langle \eta F(x_t), y \rangle + D_h(y, x_t) \} && \text{(7a)}\\
x_{t+1} &= \arg\min_{y \in \mathcal{X}} \; \{ \langle \eta F(\omega_t), y \rangle + D_h(y, x_t) \} , && \text{(7b)}
\end{aligned}
$$

Proposition 2: Let $\alpha = 1$, $\beta = 0$, $S = \nabla h - \eta F$, and $\Phi = \eta F$. Then, (5) recovers (7), and convergence follows from Corollary 1.

Proof: There exists a sufficiently small η s.t. $S = \nabla h - \eta F$ is strongly monotone. Then, $T = (\nabla h + \mathcal{N}_{\mathcal{X}})^{-1} \circ (\nabla h - \eta F) = \nabla h^* \circ (\nabla h - \eta F)$ by the fact that $(\nabla h + \mathcal{N}_{\mathcal{X}})^{-1} = \nabla h^*$. All C1–C4 and C2$^+$ are satisfied.

The first-order optimality conditions for (7) are:

$$
\begin{aligned}
\langle \nabla h(\omega_t) - \nabla h(x_t) + \eta F(x_t), x - \omega_t \rangle \geq 0, \; \forall x \quad \text{(8a)}\\
\langle \nabla h(x_{t+1}) - \nabla h(x_t) + \eta F(\omega_t), x - x_{t+1} \rangle \geq 0, \; \forall x. \quad \text{(8b)}
\end{aligned}
$$

From (8), we identify

$$
\begin{aligned}
\omega_t &= \left((\nabla h + \mathcal{N}_{\mathcal{X}})^{-1} \circ (\nabla h - \eta F) \right) (x_t) = T(x_t) && \text{(9a)}\\
x_{t+1} &= \nabla h^* (\nabla h(x_t) - \eta F(\omega_t)). && \text{(9b)}
\end{aligned}
$$

It remains to show $(S \circ T - S)(x_t) = -\eta F(\omega_t)$, so that (5) reduces to (9b): $S(T(x_t)) - S(x_t) = S(\omega_t) - S(x_t) = \nabla h(\omega_t) - \eta F(\omega_t) - (\nabla h - \eta F)(x_t) = -\eta F(\omega_t)$. ∎

The recent EG+ variant [5] generalizes EG by using different step sizes $\eta_1 > 0$ and $\eta_2 > 0$ in (7a) and (7b), respectively. TMD recovers EG+ with $\alpha = \frac{\eta_2}{\eta_1}$, stated as follows.

Corollary 3: Let $\alpha = \frac{\eta_2}{\eta_1}$, $\beta = 0$, $S = \nabla h - \eta_1 F$, and $\Phi = \eta_1 F$. Then, (5) recovers EG+, and convergence follows from Corollary 1.

[3] While discretization introduces a step size δ, this variable can be absorbed w.l.o.g. by redefining the coefficients. Thus, we omit δ in (5). Convergence follows from the continuous-time analysis under consistent discretization; a rigorous discrete-time analysis is left for future work.

Remark 1 (Weak Minty variational inequality [5]): Reference [5] also studies an L-Lipschitz F with the following property:

$$\exists\, x^* \in \mathrm{SOL}(\mathcal{X}, F) \text{ s.t. } \langle F(x), x - x^* \rangle \geq -\rho \|F(x)\|^2, \quad (10)$$

which slightly relaxes the VS requirement. It proves stability for $\rho \in [0, \frac{1}{4L})$ under the 2-norm. TMD recovers this since C2$^-$ is satisfied under (10). With the selection in Corollary 3 and h specialized to Euclidean norm, we have $T(x) - x = -\eta_1 F(x)$. Thus, the condition (4) reduces to the positivity of $\|T(x) - x\|^2 - \rho\eta_1 \|F(T(x))\|^2$. If we bound the second term via Young's inequality, then we have the strict positivity for $\rho < \frac{1}{4L}$ and $x \notin \mathrm{SOL}(\mathcal{X}, F)$.

Up to this point, Φ has appeared intrinsically tied to F (with a scaling factor). The TMD framework, however, admits a broader class of design choices for Φ. We now explore two such configurations: the first recovers several splitting methods, and the second recovers the Brown-von Neumann-Nash (BNN) dynamics, which establishes a bridge between evolutionary game dynamics and optimization algorithms.

3) Splitting methods: In splitting methods, F is decomposed as $A + B$, where A and B possess individually favorable properties. These algorithms commonly assume A and B are maximal monotone[4], and rely heavily on the resolvent operator[4] $J_X = (I + X)^{-1}$, where I is the identity map $I(x) = x$. Two well-known instances are:

(i) Douglas-Rachford splitting (DR) [6, Chapter 12.4]:

$$x_{t+1} = (J_{\eta A} \circ (2J_{\eta B} - I) + I - J_{\eta B})(x_t) := T_{\mathrm{DR}}(x_t). \quad (11)$$

(ii) Forward-Backward splitting[5] (FB) [6, Chapter 12.4], [8]:

$$x_{t+1} = J_{\eta A}\left((I - \eta B)(x_t)\right) := T_{\mathrm{FB}}(x_t). \quad (12)$$

Since the constraints can be absorbed into A or B via indicator functions, we take $\mathcal{X} = \mathbb{R}^n$ in this subsection without loss of generality, and set $\nabla h = I$ for simplicity. A notable distinction from the previous subsections is that Φ is no longer tied to F; instead, we set $\Phi = T_{\mathrm{op}}^{-1} - I$, where $T_{\mathrm{op}} \in \{T_{\mathrm{DR}}, T_{\mathrm{FB}}\}$. By this selection, T_{op} serves as the resolvent of Φ.

Proposition 3: Let $\alpha = 1$, $\beta = 0$, $S = I$, and $\Phi = T_{\mathrm{DR}}^{-1} - I$. Then, (5) recovers (11), and convergence follows from Corollary 1.

Proof: We have $T = (S + \Phi)^{-1} \circ S = T_{\mathrm{DR}}$. Since T_{DR} is firmly nonexpansive[4], Φ is (maximal) monotone, satisfying C2$^+$. To verify C3, note that the solutions correspond to the fixed points of T_{DR}, which align with the fixed points of Φ. Substituting the selections into (5) yields (11). ∎

Proposition 4: Let $\alpha = 1$, $\beta = 0$, $S = I$, and $\Phi = T_{\mathrm{FB}}^{-1} - I$. Then, (5) recovers (12), and convergence follows from Corollary 2.

Proof: The argument follows that of Proposition 3, except that T_{FB} is not firmly nonexpansive, so C2$^+$ no longer holds and we instead appeal to the relaxed condition C2$^-$.

Let $y = T_{\mathrm{FB}}(x)$, then by definition we have $y + \eta A(y) = x - \eta B(x)$ and $\Phi(y) = x - y = \eta A(y) + \eta B(x)$. Thus, condition (4) requires positivity of

$$\langle \Phi(y), y - x^* \rangle + \|y - x\|^2 \quad (13a)$$
$$\geq \eta\bar{\sigma}\|y - x^*\|^2 - \eta\bar{L}\|y - x\|\|y - x^*\| + \|y - x\|^2, \quad (13b)$$

for any fixed points x^* of T_{FB}, where (13b) is by $\bar{\sigma}$-strongly monotone $A + B$, Cauchy-Schwarz, and $\bar{L}$-Lipschitz B. For $\eta < \frac{4\bar{\sigma}}{\bar{L}^2}$, we have (13) positive unless $x = y = x^*$. ∎

4) Brown-von Neumann-Nash (BNN) dynamics [9]: The BNN dynamics from evolutionary game theory are capable of solving monotone games [9, Chapter 7.2]. They evolve on the standard simplex $\mathcal{X} = \Delta^n$. Using pre-superscripts to index vector entries, i.e., $x = [{}^1x, \ldots, {}^nx]^\top \in \Delta^n$, we write the BNN dynamics as, $i = 1, \ldots, n$,

$${}^i\dot{x}(t) = \left[-{}^i\hat{F}(x(t))\right]_+ - {}^ix(t) \sum_{k=1}^{n} \left[-{}^k\hat{F}(x(t))\right]_+, \quad (14)$$

where $\hat{F} = F - \bar{F}\mathbf{1}$, with $\bar{F}(x) = x^\top F(x)$ the weighted average of F, and $[\cdot]_+$ denotes the entrywise positive part.

In the evolutionary game literature, $[\hat{P}]_+$ is called the *excess payoff* of P, and the normalized excess payoff is defined as $[\hat{P}]_+^{\mathrm{nor}} := \frac{[\hat{P}(x)]_+}{x}$, where the division is elementwise. The key idea here is to use $[-\hat{F}(x)]_+^{\mathrm{nor}}$ to generate a target point, where the minus sign comes from F being a cost. Starting from the current x, map to the dual space via ∇h, shift by $\eta[-\hat{F}(x)]_+^{\mathrm{nor}}$, then map back:

$$T(x) = \nabla h^*\left(\nabla h(x) + \eta[-\hat{F}(x)]_+^{\mathrm{nor}}\right). \quad (15)$$

With h chosen as the negative entropy function, which is suitable for the simplex domain, (15) admits a closed form, and the TMD dynamics reduces to (14).

Lemma 2: Let $\alpha = \frac{1}{\eta}$, $\beta = 0$, $S = \nabla h$ with h the negative entropy, and T as in (15). Then, (2a) reduces to (14) and

$$T(x) = x \oplus \exp\left(\eta[-\hat{F}(x)]_+^{\mathrm{nor}}\right) \in \mathrm{int}(\Delta^n), \quad (16)$$

where $a \oplus b = \left[\frac{{}^1a\,{}^1b}{\sum_k {}^ka\,{}^kb}, \ldots, \frac{{}^na\,{}^nb}{\sum_k {}^ka\,{}^kb}\right]^\top$ for $a, b \in \mathbb{R}^n_{>0}$.[6]

Proof: From (15), we have $(S \circ T - S)(x) = \eta[-\hat{F}(x)]_+^{\mathrm{nor}}$. For $h(x) = \sum_{i=1}^n ({}^ix)\log({}^ix)$, we get $\nabla h(x) = \log x + \mathbf{1}$, $\nabla h^*(z) = \frac{\exp(z)}{\sum_{k=1}^n \exp({}^kz)}$, and $\nabla^2 h^*(\nabla h(x)) = \mathrm{diag}(x) - xx^\top$, where log and exp are applied entrywise and diag is the diagonal matrix. Differentiating $x = \nabla h^*(z)$ gives

$$\dot{x}(t) = \left(\mathrm{diag}(x) - xx^T\right)[-\hat{F}(x)]_+^{\mathrm{nor}}, \quad (17)$$

which in entrywise form yields (14). The closed form (16) follows by direct calculations. ∎

While most preceding subsections recover algorithms via discrete-time reductions, the following subsections—like BNN—operate in continuous-time, where TMD's dynamics (2) apply directly.

[4]Refer to [6, Chapter 12.3 & Chapter 12.4] for more details.

[5]Unlike DR, which only requires A and B maximal monotone, FB may require additional conditions, e.g., B Lipschitz and $A+B$ strongly monotone.

[6]The operator $\oplus$ corresponds to the vector addition in Aitchison geometry.

5) Forward-Backward-Forward dynamics [10], [11]:

$$y(t) = P_{\mathcal{X}}(x(t) - \eta F(x(t))) \tag{18a}$$

$$\dot{x}(t) = y(t) + \eta(F(x(t)) - F(y(t))) - x(t) \tag{18b}$$

Proposition 5: Let $\alpha = 1$, $\beta = 0$, $S = I - \eta F$, and $\Phi = \eta F$. Then, (2) reduces to (18), and convergence follows from Corollary 1.

Proof: With the selections, $T = (I + \mathcal{N}_{\mathcal{X}})^{-1} \circ (I - \eta F)$. Since $(I + \mathcal{N}_{\mathcal{X}})^{-1}$ is the projection onto $\mathcal{X}$, we get $y(t) = T(x(t))$. Then, (18b) is equivalent to $\dot{x}(t) = (S \circ T - S)(x(t))$. ■

To end this section, we demonstrate how we leverage the TMD framework to improve the discounted mirror descent by correcting a structural limitation in its equilibrium behavior.

6) Discounted Mirror Descent (DMD): The DMD reads

$$\dot{z}(t) = \gamma(-F(x(t)) - z(t)), \quad x(t) = \nabla h^*(z(t)), \tag{19}$$

with $\gamma > 0$. As noted in [12], the equilibrium of (19) does not lie in $\mathrm{SOL}(\mathcal{X}, F)$, but in a perturbed solution set. Our approach is to precondition $-F$ via the target correction mechanism so the equilibria are recalibrated to match $\mathrm{SOL}(\mathcal{X}, F)$.

Proposition 6: Consider the following two cases:

1) $\alpha = \gamma$, $\beta = 0$, $S = \nabla h$, $\Phi = \eta F$
2) $\alpha = \beta = \gamma$, $S = \nabla h - \eta F$, $\Phi = \eta F$.

In either case, (2) reduces to the following calibrated DMD:

$$\dot{z}(t) = \gamma\Big(\tilde{F}(x(t)) - z(t)\Big), \quad x(t) = \nabla h^*(z(t)), \tag{20}$$

where $\tilde{F} = S \circ T$, and convergence follows from Corollary 1.

Proof: Since $\alpha S + \beta \Phi = \nabla h = z$ in both cases, (2) reduces to (20). The conditions are verified as in Proposition 1 and Proposition 2 for Case 1 and Case 2, respectively. ■

Remark 2: The TMD dynamics (2a) admits a natural extension by incorporating a design input $u(t)$:

$$\dot{z}(t) = \alpha(S \circ T - S)(x(t)) - \beta \Phi(x(t)) + u(t). \tag{21}$$

Reference [12] introduces the second-order mirror descent (MD2) to improve DMD, using an equilibrium-independent passivity (EIP) argument [16]. Since (21) is output strictly EIP from input u to output x with the storage function $D_h(\bar{x}, x)$, the same technique applies here and can yield a higher-order TMD: with $\gamma_1, \gamma_2 > 0$,

$$\dot{z}(t) = \alpha(S \circ T - S)(x(t)) - \beta \Phi(t) - \gamma_1(x(t) - \xi(t)) \tag{22a}$$

$$\dot{\xi}(t) = \gamma_2(x(t) - \xi(t)). \tag{22b}$$

Setting $\alpha = 0$ recovers MD2. Moreover, while MD2 requires interior solutions, (22) is free from this restriction.

IV. Geometric Ensembles

In this section, we demonstrate the ability of TMD to run a *geometric ensemble*: multiple algorithms solve the same problem using distinct mirror maps, while their dual updates are driven by a shared ensemble state. Formally, consider N TMD instances $\mathcal{A}_i$, each equipped with a distinct mirror map $\nabla(h^i)^*$, primal state x^i, and dual state z^i, for $i = 1, \ldots, N$. Then, the proposed ensemble TMD dynamics is

$$\dot{z}^i(t) = \alpha(S \circ T - S)(x^{\mathrm{en}}(t)) - \beta \Phi(x^{\mathrm{en}}(t)) \tag{23a}$$

$$x^i(t) = \nabla(h^i)^*(z^i(t)), \quad z^i(0) \in \mathbb{R}^n, \tag{23b}$$

$$x^{\mathrm{en}}(t) = \frac{1}{N} \sum\nolimits_{k=1}^{N} x^k(t). \tag{23c}$$

For (23) to be well-posed, the mirror maps $\nabla(h^i)^*$ must share a common dual domain so that (23b) is well-defined.

The following theorem shows that the ensemble state x^{en} reduces to a single TMD instance with a synthesized mirror map—hence the name *geometric ensemble*. This is made possible by the decoupling of ∇h^* from S, T, and Φ in TMD, which allows each $\mathcal{A}_i$ to adopt a distinct mirror map without altering the shared dual update. The synthesized mirror map is given by the infimal convolution, defined as follows.

Definition 4 (Infimal Convolution [7, Chapter 12]): Let $\mathbb{V}$ be a vector space and $f, g : \mathbb{V} \to \bar{\mathbb{R}} = \mathbb{R} \cup \{+\infty\}$. Their infimal convolution, denoted by $f \,\square\, g : \mathbb{V} \to \bar{\mathbb{R}}$, is

$$(f \,\square\, g)(v) = \inf\{f(a) + g(b) : a, b \in \mathbb{V}, a + b = v\}. \tag{24}$$

Lemma 3 ([7, Chapter 12–15]): Let f, g be proper convex functions. Then, $f \,\square\, g$ is convex and $(f \,\square\, g)^* = f^* + g^*$.

Theorem 2: If each $\mathcal{A}^i$ has initial dual state $z^i(0) \in \mathbb{R}^n$, then the ensemble state x^{en} in (23) satisfies

$$\dot{z}(t) = \alpha(S \circ T - S)(x^{\mathrm{en}}(t)) - \beta \Phi(x^{\mathrm{en}}(t)) \tag{25a}$$

$$x^{\mathrm{en}}(t) = \nabla(h^{\mathrm{en}})^*(z(t)), \text{ with } z(0) = 0, \tag{25b}$$

where $h^{\mathrm{en}} = \bar{h}^i \,\square\, \ldots \,\square\, \bar{h}^N$ with $\bar{h}^i(x) = \frac{1}{N} h^i(Nx) - \langle z^i(0), x \rangle$.

Proof: Since (25a) is common to all agents, $z^i(t) - z^j(t)$ depends only on initial values and forms a fixed pattern in the dual space. Thus, we can describe the dual update consistently as $\dot{z} = \alpha(S \circ T - S)(x^{\mathrm{en}}) - \beta \Phi(x^{\mathrm{en}})$ with $z^i(t) = z(t) + z^i(0)$. Note that, by definition, we get

$$(\bar{h}^i)^*(z) = \sup_x \left\{ \langle z, x \rangle - \frac{1}{N} h^i(Nx) + \langle z^i(0), x \rangle \right\} \tag{26a}$$

$$= \frac{1}{N} (h^i)^*(z + z^i(0)), \tag{26b}$$

Then, (23b) and (23c) together yield

$$(h^{\mathrm{en}})^*(z) = \sum_{k=1}^{N} (\bar{h}^k)^*(z) = \left(\bar{h}^1 \,\square\, \ldots \,\square\, \bar{h}^N\right)^*(z), \tag{27}$$

where (27) is by (26) and Lemma 3. ■

Theorem 2 states that the ensemble behaves as a single TMD with synthesized mirror map $\nabla(h^{\mathrm{en}})^*$, and thus enjoys the convergence results from Theorem 1 or Corollary 1. We demonstrate the relative robust convergence of ensemble TMD compared to the individual TMD in Section V. The tradeoff is maintaining the N primal states. In practice, updates can be made in parallel, keeping the computational overhead modest.

V. Simulations

We present two simulations validating the TMD framework.

TMD vs. Mirror Descent: We consider $\mathrm{VI}(\mathbb{R}^2, F)$ with merely monotone $F(x) = A(x - x^*)$, where $A = \begin{bmatrix} 0 & -3 \\ 3 & 0 \end{bmatrix}$ and $x^* = [2, 2]^\top$. We compare mirror descent against three TMD instances, all with $\Phi = \eta F$ and $T = (S + \Phi)^{-1} \circ S$, but distinct ∇h and S: (i) $\nabla h(x) = x$, $S(x) = x + x^3$

(component-wise); (ii) $\nabla h(x) = Hx$ ($H \succ 0$), $S(x) = x$; (iii) $\nabla h(x) = -1/x$ (log-barrier), $S(x) = x$. As shown in Figure 2, mirror descent diverges under forward Euler discretization while all three TMD instances converge, confirming Corollary 1. The diversity across instances demonstrates that TMD accommodates a broad class of design choices beyond the specific selections recovering PPA or EG.

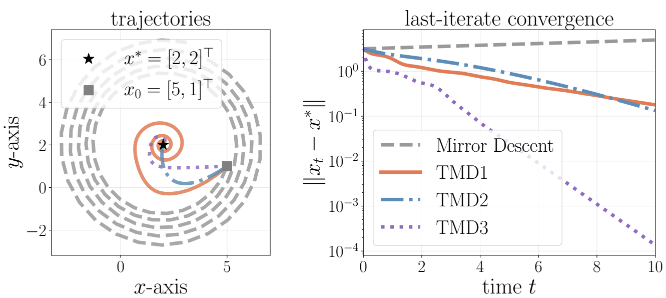

Fig. 2. Mirror descent diverges while all TMD instances converge.

Geometric Ensemble: As in mirror descent, the selection of ∇h requires prior knowledge of the problem geometry. Ensemble TMD partially addresses this since it equivalently synthesizes multiple candidate mirror maps, as stated in Theorem 2. We consider three distinct mirror maps, $\nabla h^1, \nabla h^2, \nabla h^3$, under three test cases: $\mathrm{VI}(\mathbb{R}^3, F)$ where F is constructed based on each of the three mirror maps in turn. As shown in Figure 3, in each case, the mirror map whose geometry aligns with F converges fastest, while the remaining two are slow. Ensemble TMD achieves consistent convergence across all three cases.

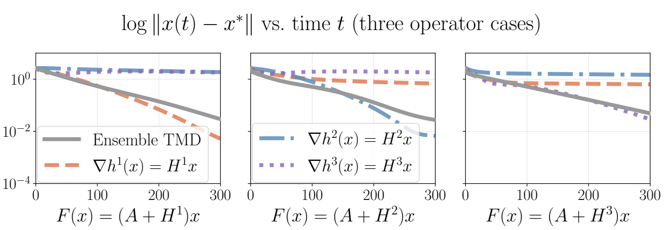

Fig. 3. The matched $\boldsymbol{h^i}$ converges fastest under the corresponding case, while ensemble TMD achieves consistent convergence across all three cases.

VI. Conclusions

In this paper, we proposed Target Mirror Descent (TMD), a unified framework for solving monotone variational inequalities via a target point correction mechanism in the dual update. We established rigorous convergence guarantees and demonstrated TMD's unification of several landmark algorithms. Moreover, TMD decouples the mirror map from the target determination, thereby enabling *geometric ensembles*. Numerical simulations confirmed the theoretical guarantees and consistent convergence of geometric ensembles across varying problem geometries. It would be interesting to generalize TMD to non-autonomous systems, where time-varying S and T may capture convergence acceleration techniques, and time-varying Φ non-stationary environments.

References

[1] G. M. Korpelevich, "The extragradient method for finding saddle points and other problems," *Matecon*, vol. 12, pp. 747–756, 1976.

[2] P. Mertikopoulos and M. Staudigl, "Stochastic Mirror Descent Dynamics and Their Convergence in Monotone Variational Inequalities," *Journal of Optimization Theory and Applications*, vol. 179, no. 3, pp. 838–867, 2018.

[3] J. Eckstein, "Nonlinear Proximal Point Algorithms Using Bregman Functions, with Applications to Convex Programming," *Mathematics of Operations Research*, vol. 18, no. 1, pp. 202–226, 1993.

[4] C. D. Dang and G. Lan, "On the convergence properties of non-Euclidean extragradient methods for variational inequalities with generalized monotone operators," *Computational Optimization and Applications*, vol. 60, no. 2, pp. 277–310, 2015.

[5] J. Diakonikolas, C. Daskalakis, and M. I. Jordan, "Efficient methods for structured nonconvex-nonconcave min-max optimization," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2021, pp. 2746–2754.

[6] F. Facchinei and J.-S. Pang, *Finite-Dimensional Variational Inequalities and Complementarity Problems*, ser. Springer Series in Operations Research. Springer, 2003.

[7] H. H. Bauschke and P. L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, ser. CMS Books in Mathematics. Springer International Publishing, 2017.

[8] M. N. Bùi and P. L. Combettes, "Bregman Forward-Backward Operator Splitting," *Set-Valued and Variational Analysis*, vol. 29, no. 3, pp. 583–603, 2021.

[9] W. H. Sandholm, *Population Games and Evolutionary Dynamics*, ser. Economic Learning and Social Evolution. MIT Press, 2010.

[10] R. I. Boţ, E. R. Csetnek, and P. T. Vuong, "The forward–backward–forward method from continuous and discrete perspective for pseudo-monotone variational inequalities in Hilbert spaces," *European Journal of Operational Research*, vol. 287, no. 1, pp. 49–60, 2020.

[11] P. Tseng, "A Modified Forward-Backward Splitting Method for Maximal Monotone Mappings," *SIAM Journal on Control and Optimization*, vol. 38, no. 2, pp. 431–446, 2000.

[12] B. Gao and L. Pavel, "Second-Order Mirror Descent: Convergence in Games Beyond Averaging and Discounting," *IEEE Transactions on Automatic Control*, vol. 69, no. 4, pp. 2143–2157, 2024.

[13] A. Mokhtari, A. Ozdaglar, and S. Pattathil, "A Unified Analysis of Extra-gradient and Optimistic Gradient Methods for Saddle Point Problems: Proximal Point Approach," in *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 1497–1507.

[14] P. Mertikopoulos and Z. Zhou, "Learning in games with continuous action sets and unknown payoff functions," *Mathematical Programming*, vol. 173, no. 1-2, pp. 465–507, 2019.

[15] H. H. Bauschke, J. M. Borwein, and P. L. Combettes, "Bregman Monotone Optimization Algorithms," *SIAM Journal on Control and Optimization*, vol. 42, no. 2, pp. 596–636, 2003.

[16] G. H. Hines, M. Arcak, and A. K. Packard, "Equilibrium-independent passivity: A new definition and numerical certification," *Automatica*, vol. 47, no. 9, pp. 1949–1956, 2011.