New Trends and Challenges in Stochastic Differential Games

Pierre Cardaliaguet (Paris Dauphine-PSL University)
Miryana Grigorova (University of Warwick)
Thibaut Mastrolia (UC Berkeley)
Jinniao Qiu (University of Calgary)
Jianfeng Zhang (University of Southern California)

June 23–28, 2024

1 Overview of the Field

The stochastic differential games (SDGs) are concerned with stochastic dynamics governed by N players, typically modelled as follows:

$$\dot{X}^i(t) = F^i(t, X^1, \cdots, X^N, v^1, \cdots, v^N, \dot{W}^i),\, t > 0; \quad X^i(0) = x^i, \quad i = 1, \dots, N. \quad (*)$$

Here $X^i = (X^i(t))_{t\geq 0}$ is the i-th player's state process driven by the noise process $W^i = (W^i(t))_{t\in\geq 0}$ and subject to all players' strategies $(v^1, \cdots, v^N)$. Each player has the individual objective functional $J^i(X^i, v^1, \cdots, v^N)$ to maximize or minimize over the set of his/her own admissible strategies $v^i = (v^i(t))_{t\geq 0}$ for $i = 1, \dots, N$, and under such an interdependent framework, the solution is a set of strategies $(\bar{v}^1, \cdots, \bar{v}^N)$ achieving the equilibrium. The study of SDGs dates back to the 1960s (see [22]). Over the following half-century, a wide range of methods and mathematical theories have been developed and applied across diverse fields, including engineering, operations research, biology, ecology, finance, economics, the energy industry, environmental sciences, and computer science.

Over the past decade, new and intriguing classes of SDGs have emerged, extending beyond the conventional framework. These novel SDGs, which naturally arise from various applications, are notably nonlinear and non-Markovian. Recently, a diverse array of approaches has been proposed and developed to tackle these challenges, including rough path techniques, path-dependent analysis, mean-field limit methods, stochastic viscosity solutions, and mean-field forward-backward stochastic (partial) differential equations, among others. Simultaneously, various numerical methods have been introduced for approximating solutions.

This workshop brought together both pure and applied mathematicians working on different aspects of SDGs or seeking resolutions through various methodological frameworks. Participants, ranging from senior mathematicians to early career researchers and doctoral students, gathered to discuss recent achievements and challenges in the field, potentially illuminating new research trends.

2 Recent Developments and Open Problems

2.1 Recent Developments

During the workshop, the presentations and discussions highlighted some recent developments and emerging trends in Stochastic Differential Games (SDGs), which can be broadly categorized into the following five perspectives.

(I) Some new frameworks and novel approaches for SDGs were presented and discussed.

- The α-potential games is raised as a a new paradigm for N-player games. The analytical criteria for any game to be an α-potential game was presented, and several important classes of Markov α-potential games were discussed. The talks also covered detailed analysis for games with mean-field interactions, distributed games, and the crowd aversion games, in which α is shown to depend on the number of players, the admissible policies, and the cost structure. The changes of α from open-loop to closed-loop settings were also investigated. We refer to the work [19] by Xin Guo, Xinyu Li, and Yufei Zhang, and references therein for more details.
- Using the theory of quantum filtering, a framework for quantum dynamic games was proposed, where players choose their strategies in real time. Quantum filtering enables the reduction of quantum dynamic games to stochastic differential games with diffusive or jump-type noises on complex Riemannian manifolds, which are generally infinite-dimensional. Quantum mean-field games have been reformulated and studied as a specific class of controlled infinite-dimensional McKean-Vlasov diffusions on complex projective spaces. This development has introduced a variety of new problems and challenges, including the study of a new class of nonlinear stochastic differential equations in Banach spaces with singular coefficients, which arise as the law of large number limits of stochastic quantum master equations for mixed states. Some fundamental features of this emerging theory, along with several solvable examples, were also demonstrated. For further references, see [27, 28] by Vassili N. Kolokoltsov.
- The recent work by Cvitanić, Possamaï, and Touzi (2018) [12] presents a general approach to continuous-time principal-agent problems using dynamic programming and second-order backward stochastic differential equations (2BSDEs). In contrast, an alternative formulation of the principal-agent problem was presented at the workshop, which can be solved using more straightforward techniques, relying solely on the theory of BSDEs; see [10] where the authors believe this more straightforward approach to solve general continuous-time principal-agent problems with volatility control will facilitate the dissemination of these problems across many fields, and its extension to more intricate principal-agent problems, for example with many (or even a continuum of) agents, or to more general, non-necessarily continuous, output processes. Furthermore, a general target approach was proposed to reformulate any continuous-time stochastic Stackelberg differential game with closed-loop strategies as a single-level optimization problem with target constraints, and the aim was to characterize the Stackelberg equilibrium when players use 'closed-loop strategies,' meaning their decisions rely solely on the historical information of the output process, specifically excluding any direct dependence on the underlying driving noise, which was often unobservable in real-world applications. The first step was to demonstrate that by treating the second-order backward stochastic differential equation associated with the follower's continuation utility as a controlled state variable for the leader, the leader's unconventional optimization problem can be reformulated as a more conventional stochastic control problem with stochastic target constraints. Subsequently, by adapting the methodologies developed by Soner and Touzi or Bouchard, Élie, and Imbert, the optimal strategies and the corresponding value of the Stackelberg equilibrium could be characterized through the solution of a well-defined system of Hamilton–Jacobi–Bellman equations; see [21] by Hernández, Santibáñez, Hubert and Possamaï.
- A new approach to solving mean-field games was introduced by Dumitrescu Dumitrescu, Leutscher and Tankov [15]. This approach is based on reformulating stochastic control problems as linear programming problems. The review covered the linear programming approach for mean field games (MFG) of control and optimal stopping, with a discussion on numerical algorithms and

applications to energy transition. The focus was on a specific model related to the long-term dynamics of the electric market structure under potential uncertainty, driven by the progressive replacement of conventional power generation with renewable energy sources.

- A classical model for ticket pricing, introduced by Gallego and van Ryzin (1994), assumes a single seller and an infinite time horizon. This model has been widely utilized by airlines, among others. However, extending the model to more realistic settings—where multiple sellers with finite inventories compete over a finite time horizon—is mathematically and computationally more complex. To address these challenges, Ronnie Sircar proposed and studied a new class of stochastic games involving intensity control for ticket pricing.

(II) Some new results on associated Master equations and Hamilton-Jacobi-Bellman-Isaacs equations were presented.

- When the concerned dynamics allows for path/history-dependent coefficients, e.g., F^i and/or J^i depend on the path of state X^i and/or the paths of the strategies $(v^1, \cdots, v^N)$, the resulting SDG is beyond the conventional Markovian framework; for the zero-sum SDGs, a class of path-dependent Hamilton-Jacobi-Bellman-Isaacs (HJBI) equations are proposed and studied; some new results on the existence and uniqueness of viscosity solution to such kind of fully nonlinear second order path-dependent partial differential equations were presented (see Tang and Zhou [37]). When the coefficients are generally random, e.g., F^i and/or J^i depend on the path of the Brownian motion W^i in a measurable manner or allow for randomness adapted to the filtration generated by W^i, there follows the non-Markovian zero-sum SDG from which the stochastic HJBI equation arises naturally for the characterization of the value function (see Qiu and Zhang [36]) and a viscosity solution theory on such stochastic HJBI equations was presented in the workshop. In addition, an Itô's formula for C^1-functionals of flows of conditional marginal distributions of continuous semimartingales was introduced (see Bouchard, Tan and Wang [3]) and a class of McKean-Vlasov optimal control problems and associated (viscosity) solution of the associated HJB master equation were studied, where the verification theorem only requires C^1-regularity of its value function. A comprehensive viscosity solution theory on a class of second order Hamilton-Jacobi-Bellman (HJB) equations in the Wasserstein space was also presented (refer to Cheung, Tai and Qiu [9]) where the associated mean field control problems allow for common noises. The wellposedness of weak solutions for a class of parabolic partial differential equations in the Wasserstein space was studied and the connections between these partial differential equations, mean-field control and filtering problems were established in Bayraktar, Ekren and Zhang [1].
- Some novel results on master equations for mean-field games and controls are also presented. Mean-field game refers to a class of SDGs allowing for coefficients F^i and J^i depending on the distribution or conditional distributions of the state process X^i and it is the limit of a sequence of $N-$player SDGs as N tends to infinity. In the mean-field game theory certain classes of forward-backward stochastic partial differential equations (FBSPDEs) arise as the mean-field game system with common noise; the former is a forward stochastic Kolmogorov equation describing the evolution of the conditional distributions of the states of the players given the common noise, while the latter is the stochastic Hamilton-Jacobi-Bellman (HJB) equation characterizing the value function of the optimization problem when the flow of conditional distributions is given (see Cardaliaguet, Delarue, Lasry and Lions [4] and Carmona and Delarue [5]). Meanwhile, the master equation, a nonlocal second-order nonlinear partial differential equation defined on the space of probability measures, is proposed and studied for the mean-field games. The achieved results on master equations were presented and discussed. Particularly, for mean field Stackelberg games with a leader and a large number of followers, the dynamic programming approach and associated HJB equations (also called master equations) were studied and discussed in a detailed way within a linear quadratic framework; see Yang and Huang [38]. The well-posedness of the master equation, with results in the case in which the dynamics of the common noise depend on the decisions of the players, were also analyzed by Charles Bertucci and his coauthors, with systematic study of common noise in MFG with the common randomness affects players through

an exterior factors. A collection of recent results on mean field games and the corresponding master equations governed by non-separable Hamiltonians and potentially degenerate idiosyncratic noise were presented, where notion of displacement monotonicity ensures global in time well-posedness theories and give particular stabilization effects when one considers long time behavior of the solutions; see Gangbo and Mészáros [17], and Mészáros and Mou [33] for instance. Minimal solutions of master equations for extended mean field games were proposed and studied in Mou and Zhang [34].

(III) Various numerical methods have been proposed for the kind of equations above. Due to the nonlinear, non-convex, non-local, non-Markovian, and/or possibly degenerate properties of the involved equations, one needs to deal with the arising large time-complexity and high-dimensionality. Some novel numerical methods were discussed and presented, which fostered the developments and applications of numerical solutions to SDGs.

- Deep backward and Galerkin methods were investigated for learning finite state master equations in Cohen, Laurière and Zell [11]. The first method we propose relies on backward induction while the second one directly tackles the master equation without discretizing time. It is proved that: there exist neural networks that make the loss functions of the algorithms arbitrarily small and, conversely, if the losses are small, then the neural networks are good approximations of the master equation solution.
- Modeling multi-agent systems on networks is a fundamental challenge in a wide variety of disciplines. The weight matrix of the network and the interaction kernel were jointly inferred in Lang, Wang, Lu and Maggioni [30], which determine respectively which agents interact with which others and the rules of such interactions from data consisting of multiple trajectories. Estimator were proposed, which leads naturally to a non-convex optimization problem, and two approaches were investigated for its solution: one is based on the alternating least squares algorithm; another is based on a new algorithm named operator regression with alternating least squares.
- Continues-time reinforcement learning method was studied and presented for optimal stopping problems as in Dianetti, Ferrari and Xu [13]. Also,the mean field games (MFGs) were used to investigate approximations of N-player games with heterogeneous symmetrically continuous closed-loop actions. Centered around the notion of regret, non-asymptotic analysis was conducted on the approximation capability of MFGs from the perspective of state-action distributions without requiring the uniqueness of equilibria. See Cheng and Jaimungal [8].
- Using dynamic programming and a forward-backward algorithm, the mean field equilibria of multi period models was constructed as concatenation of equilibria of one-step games and results on convergence of discrete time games to continuous time counterparts were also presented by Ludovic Tangpi.
- Policy improvement algorithms for reinforcement learning were studied for continuous-time entropy-regularized stochastic control problems in Ma, Wang and Zhang [32], where uniform convergence both for the iterative value functions and for their derivatives was proved.
- Convergence results of the simulation of the density solution to the McKean-Vlasov equation were proved when the measure variable is in the drift in Hoffmann and Liu [23]. Such a statistical method builds upon adaptive nonparametric results. A generalised Bernstein inequality was obtained for Euler schemes with interacting particles and obtain sharp deviation inequalities for the estimated classical solution.
- Inspired by swarm intelligence, novel consensus-based optimization approaches were proposed and studied in [2, 7, 24] as presented by Hui Huang for nonconvex-nonconcave min-max problems, utilizing a multi-particle, metaheuristic, derivative-free method capable of provably finding global solutions. Also, a class of self-interacting approximations were proposed and investigated for the invariant measure of McKean-Vlasov long time limit in Du, Ren, Suciu and Wang [14], where if the McKean-Vlasov diffusion is the gradient flow of a convex mean-field potential functional, the self-interacting process is shown to exponentially converge towards its unique invariant measure close to that of the McKean-Vlasov diffusion.

(IV) Emerging trends and challenges in the application of stochastic differential games across various research fields were also discussed.

- Some time-inconsistent stopping problems for a continuous-time Markov chain under finite time horizon with non-exponential discounting were considered as games and Zhou Zhou and his coauthors in [25] proposed a notion of equilibrium called almost strong equilibrium (ASE), which is a weak equilibrium and satisfies the condition of strong equilibria except at the boundary points of the associated stopping region. An iteration procedure leading to an ASE was provided and furthermore, this ASE was shown to be the unique ASE among all regular stopping policies under finite horizon.
- Bin Zou and his coauthors in [31] used a two-layer stochastic game model to study reinsurance contracting and competition in a market with one insurer and two competing reinsurers. The insurer negotiates with both reinsurers simultaneously for proportional reinsurance contracts that are priced using the variance premium principle; the reinsurance contracting between the insurer and each reinsurer is modeled as a Stackelberg game. The two reinsurers compete for business from the insurer and optimize the so-called relative performance, instead of their own surplus; the competition game between the two reinsurers is settled by a non-cooperative Nash game. A sufficient and necessary condition for unique existence of the equllibrium and sensitivity results were provided.
- In the Stefan problem with Gibbs-Thomson law, the characterization of the correct times, directions and sizes of the jumps of the associated free boundary along the time variable in a general multidimensional setting had remained an open question, until recently, Sergey Nadtochiy and M. Shkolnikov and Y. Guo (see [20] for reference) derived a separate (hyperbolic) partial differential equation — referred to as the cascade equation — whose solutions describe the jumps of the solutions to the Stefan problem without any symmetry assumptions. It turns out that a solution of the cascade equation corresponds to a maximal element of the set of all equilibria in a family of (first-order local) mean field games.
- Both N-player and mean-field games arising from optimal portfolio liquidation were analyzed, where players were restricted from altering the direction of trading. Specifically, players with an initially short position in stocks were only permitted to buy, while those with an initially long position could only sell. Under appropriate conditions on the model parameters, it was demonstrated that these games could be reduced to timing games, where players must determine the optimal moments for market entry and exit. The equilibrium entry and exit times were identified, and it was shown that the equilibrium mean-trading rates can be characterized by solutions to a highly non-linear, higher-order integral equation with an endogenous terminal condition. The study addressed both the existence of a unique solution to this integral equation and the existence of a unique equilibrium in both the mean-field and N-player games. For further details, refer to the work by Ulrich Horst and his coauthors [16].
- An extension of the renowned Sannikov Principal-Agent problem to the multi-agent scenario was examined, with a focus on the corresponding mean-field model involving a continuum of agents. It was demonstrated that the Principal's problem can be reduced to a mixed control-and-stopping mean-field problem, and a semi-explicit solution to the first-best contracting problem was derived. For further details, refer to the joint work by Mehdi Talbi, Thibaut Mastrolia, and Nizar Touzi.

(V) The workshop also addressed topics pertinent to both the theoretical and practical advancements in SDGs.

- A related optimal transport problem with backward martingale constraint was discussed by Mihai Sirbu [29]. When the objective function is given by the scalar product of a pseudo-Euclidean space $\mathcal{S}$, it is shown that the supremums over maps and plans coincide, provided that the law ν of the input random variable Y is atomless. An optimal map X exists if ν does not charge any c-c surface (the graph of a difference of convex functions) with strictly positive normal vectors in the sense of the $\mathcal{S}$-space. The optimal map X is unique if ν does not charge c-c surfaces with nonnegative normal vectors in the $\mathcal{S}$-space. As an application, sharp conditions were derived for the existence and uniqueness of equilibrium in a multi-asset version of the model with insider from Rochet and Vila. In the linear-Gaussian case, Kyle's lambda, the sensitivity of price to trading volume, is characterized as the unique positive solution of a non-symmetric algebraic Riccati equation.

- Two studies focused on energy markets were presented. Mike Ludkovski addressed the daily operation of hybrid energy resources that combine a renewable generator with a battery energy storage system (BESS).A stochastic control formulation for the optimal dispatch of BESS was given to maximize the reliability of the hybrid asset concerning a given day-ahead dispatch target or forecast. A machine-learning algorithm based on Gaussian Process regression was introduced to efficiently derive the dynamic feedback control map, and several numerical case studies based on realistic asset simulations were showcased. For further details, see [6]. On the other hand, Yang Yang [35] investigated price-storage dynamics in natural gas markets. A novel stochastic path-dependent volatility model was proposed, with path dependence introduced in both price volatility and storage increments. The model was calibrated for both price and storage dynamics. Additionally, the pricing problem of a specific discrete swing option was explored, and a deep learning-based method was proposed with convergence analysis.
- Adapted optimal transport has recently emerged as a valuable tool for quantifying distributional uncertainty and assessing the sensitivity of stochastic optimization problems, particularly in contexts where the flow of information over time is crucial. Ting Kam Leonard Wong [18] discussed the adapted (bicausal) optimal transport between real-valued Gaussian processes in discrete time, characterizing the optimal coupling and comparing it with the classic Bures-Wasserstein case, with some derived geometric implications.
- A a new notion of set valued PDEs was proposed and studied by Melih Iseri and Jianfeng Zhang [26]. The key component in the theory was a set valued Itô formula, characterizing the flows on the surface of the dynamic sets. In the contexts of multivariate control problems, the wellposedness of the set valued HJB equations was established, which extended the standard HJB equations in the scalar case to the multivariate case.
- Hao Xing presented some interesting results on optimal contract, habit formation, and capital structure. A continuous-time principal-agent model was examined where agent's preference exhibits habit formation over consumption. As agent's concern over the standard of living strengthens, the continuation utility was less sensitive to current wealth but more sensitive to the standard of living, leading to lower demand for risk-sharing compensation. The optimal contract had lower pay-for-performance but incentivizes agent's higher effort. In the Leland (1994) capital structure model, agent's habit formation preference combined with the optimal contract would lower firm's leverage and mitigate the debt-overhang problem.

2.2 Open Problems

During the workshop, two open-problem sessions took place on Monday and Thursday afternoons, with six speakers involved. The following is a summary of some of the open problems discussed during these sessions.

(a) In the Stefan problem, Sergey Nadtochiy and the coauthors derived a separate hyperbolic partial differential equation, known as the cascade equation, whose solutions characterize the jumps in the solutions to the Stefan problem without requiring any symmetry assumptions. Remarkably, a solution to the cascade equation corresponds to a maximal element within the set of all equilibria in a family of first-order local mean field games. The existence of a (so-called minimal) solution to the cascade equation was established. An open question was raised regarding the general well-posedness of such cascade equations.

(b) In his talk, Bin Zou introduced a two-layer stochastic game model to study reinsurance contracting and competition in a market involving one insurer and two competing reinsurers. The reinsurers compete for business from the insurer by optimizing their relative performance, rather than their own surplus. The competition between the two reinsurers is modeled as a non-cooperative Nash game. Bin Zou presented both sufficient and necessary conditions for the unique existence of the equilibrium, along with sensitivity analysis results. A natural open question is how a general stochastic differential game framework, such as mean-field games, can be applied to models in insurance markets.

(c) In Renyuan Xu's open problem session, she began by reviewing the interplay between the following areas:

- Stochastic control and machine learning theory,
- Stochastic games (mean-field games) and machine learning theory.

Given that the development and understanding of large language models and generative AI are predominantly at an empirical level, Renyuan Xu initiated the following discussions with several general open questions: Can we develop stochastic control frameworks for these emerging applications to

- better understand when and why it works?
- design new algorithms based on HJB, stochastic maximum principle and FBSDE frameworks?
- improve the algorithmic efficiency based on the property (such as regularity) of the control formulation?

(d) Hao Xing's open problem session focused on dynamic debt management. He began by reviewing the Dynamic Trade-off Theory proposed by Leland (1994), followed by an introduction to theories on leverage dynamics without commitment, as developed by Demarzo and He (2021). He then discussed extensions to the model, considering scenarios with fixed debt issuance costs, short- and long-term debt, and non-Markovian strategies. Finally, Hao Xing presented two open questions:

- How can the benefits of debt be restored?
- How can the model be brought closer to empirical data?

(e) In his open problem session, Bruno Bouchard began with the topic of Asian option pricing, where the derivatives $\partial_t \varphi$ and $\partial_{xx}\varphi$ were not well-defined for the value function φ until Dupire's derivatives on path spaces were introduced. Basic calculations using parametrization techniques led to well-defined Gâteaux (directional) derivatives. The open question that arose concerns the extension to general path-dependent models. Specifically, if the state process $(X_t)_{t\in[0,T]}$ is defined through a stochastic Volterra integral equation with, for example, power-type kernels, can similar calculations be used to derive the partial differential equations (with well-defined derivatives) for the value function $V(t, X_t) = \mathbf{E}[g(X_T)\mathcal{F}_t]$?

(f) In his talk, Alpár Mészáros presented a series of recent results on mean-field games and the corresponding master equations governed by non-separable Hamiltonians and potentially degenerate idiosyncratic noise. He demonstrated how the concept of displacement monotonicity ensures global well-posedness over time and provides particular stabilization effects when examining the long-term behavior of solutions. During his open problem session, Alpár focused on various selections of master equations (or equivalently, mean-field systems), particularly with the uniqueness or non-uniqueness of solutions, which was the primary open question discussed.

3 Presentation Highlights

There were two keynote speakers at the workshop.

The first was Xin Guo from UC Berkeley, who also opened the workshop with a talk on a novel class of α-potential games, a new paradigm for N-player games. Static potential games, introduced by Monderer and Shapley (1996), are non-cooperative games where an auxiliary function, known as a static potential function, exists. This function allows any change in a player's utility upon unilaterally deviating from their strategy to be evaluated through the change in the potential function's value. The use of a potential function is powerful as it simplifies the otherwise complex task of finding Nash equilibria in multi-agent non-cooperative games: maximizers of the potential function correspond to the game's Nash equilibria. In her talk, Xin proposed an analogous framework for dynamic N-player games called α-potential games, where the static potential function was replaced by an α-potential function. She presented analytical criteria for a game to qualify as an α-potential game and identified several key classes of Markov α-potential games. Her detailed analysis also covered games with mean-field interactions, distributed games, and crowd aversion games, where α was shown to depend on the number of players, admissible policies, and the cost structure. The transition of α from open-loop to closed-loop settings was also demonstrated. Following Xin's talk, Xinyu Li presented on Markov α-potential games, and Yufei Zhang discussed an analytical framework for potential games.

The second keynote speaker was Bruno Bouchard, who delivered a talk on a C^1-Itô's formula for flows of semimartingale distributions. He presented an Itô's formula for C^1-functionals of flows of conditional marginal distributions of continuous semimartingales, extending the C^1-Itô's formula of Gozzi and Russo (2006) to this context. This was based on the concept of weak Dirichlet processes. As a first application, a class of McKean-Vlasov optimal control problems was studied, with a derived verification theorem that requires only C^1-regularity of its value function, which is equivalent to the (viscosity) solution of the associated HJB master equation. This result was complemented by a novel duality result. An example was also provided where the required regularity was verified using only probabilistic tools.

On Tuesday afternoon of the workshop, a session was held featuring talks by four PhD students: Xinyu Li, Hang Cheung, Gaozhan Wang, and Yang Yang. Additionally, four posters were presented by early-career researchers: Xiong Wang, Junxi Zhang, Yang Yang, and Gaozhan Wang.

4 Outcome of the Meeting

The workshop successfully achieved its objectives, facilitating a vibrant exchange of ideas on recent developments in stochastic differential games (SDGs). Participants shared their latest research findings and methods, showcasing significant achievements in the proposed topics. This exchange not only highlighted the current state of the field but also set the stage for future advancements.

One of the key outcomes was the in-depth discussion of emerging trends and challenges in various methodological frameworks for addressing SDGs. These discussions helped identify the limitations of existing approaches and the potential for new techniques to overcome these challenges. The collaborative environment fostered during the workshop allowed participants to explore these issues in detail, leading to the identification of possible solutions and new research directions.

For each specific topic covered, the workshop provided a platform to critically assess the challenges posed by existing methodologies. This assessment led to fruitful discussions on potential collaborations aimed at developing innovative solutions, which are expected to drive future research in the field.

The workshop also played a pivotal role in fostering novel collaborations across different approaches and topics within the SDG framework. By bringing together experts from diverse backgrounds, the event enabled participants to explore interdisciplinary connections and identify opportunities for collaborative research that might not have been evident otherwise.

Finally, the workshop was successful in fostering interactions and collaborations among participants at different career stages and with varied expertise. This inclusive environment not only enriched the discussions but also laid the foundation for ongoing mentorship, knowledge sharing, and joint research efforts, ensuring that the field of SDGs continues to evolve in a collaborative and dynamic manner.

References

[1] Erhan Bayraktar, Ibrahim Ekren, and Xin Zhang. Prediction against a limited adversary. *Journal of Machine Learning Research*, 22(72):1–33, 2021.

[2] Giacomo Borghi, Hui Huang, and Jinniao Qiu. A particle consensus approach to solving nonconvex-nonconcave min-max problems. *arXiv preprint arXiv:2407.17373*, 2024.

[3] Bruno Bouchard, Xiaolu Tan, and Jixin Wang. A C^1-Itô's formula for flows of semimartingale distributions. *Applied Mathematics & Optimization*, 90(1):26, 2024.

[4] Pierre Cardaliaguet, François Delarue, Jean-Michel Lasry, and Pierre-Louis Lions. *The Master Equation and the Convergence Problem in Mean Field Games*, volume 201 of *Annals of Mathematics Studies*. Princeton: Princeton University Press, 2019.

[5] René Carmona and Francois Delarue. The master equation for large population equilibriums. In *Crisan D., Hambly B., Zariphopoulou T. (eds) Stochastic Analysis and Applications 2014*, volume 100 of *Springer Proceedings in Mathematics & Statistics*, pages 77–128. Springer, Cham, 2014.

[6] Tao Chen, Mike Ludkovski, and Moritz Voß. On parametric optimal execution and machine learning surrogates. *Quantitative Finance*, 24(1):15–34, 2024.

[7] Enis Chenchene, Hui Huang, and Jinniao Qiu. A consensus-based algorithm for non-convex multiplayer games. *arXiv preprint arXiv:2311.08270*, 2023.

[8] Ziteng Cheng and Sebastian Jaimungal. Mean field regret in discrete time games. *arXiv preprint arXiv:2301.06930*, 2023.

[9] Hang Cheung, Ho Man Tai, and Jinniao Qiu. Viscosity solutions of a class of second order hamilton-jacobi-bellman equations in the wasserstein space. *arXiv preprint arXiv:2312.10322*, 2023.

[10] Alessandro Chiusolo and Emma Hubert. A new approach to principal-agent problems with volatility control. *arXiv preprint arXiv:2407.09471*, 2024.

[11] Asaf Cohen, Mathieu Laurière, and Ethan Zell. Deep backward and galerkin methods for the finite state master equation. *arXiv preprint arXiv:2403.04975*, 2024.

[12] Jakša Cvitanić, Dylan Possamaï, and Nizar Touzi. Dynamic programming approach to principal–agent problems. *Finance and Stochastics*, 22:1–37, 2018.

[13] Jodi Dianetti, Giorgio Ferrari, and Renyuan Xu. Exploratory optimal stopping: A singular control formulation. *arXiv preprint arXiv:2408.09335*, 2024.

[14] Kai Du, Zhenjie Ren, Florin Suciu, and Songbo Wang. Self-interacting approximation to mckean-vlasov long-time limit: a markov chain monte carlo method. *arXiv preprint arXiv:2311.11428*, 2023.

[15] Roxana Dumitrescu, Marcos Leutscher, and Peter Tankov. Control and optimal stopping mean field games: a linear programming approach. *Electronic Journal of Probability*, 26:1–49, 2021.

[16] Guanxing Fu, Paul P Hager, and Ulrich Horst. Mean-field liquidation games with market drop-out. *Mathematical Finance*.

[17] Wilfrid Gangbo and Alpár R Mészáros. Global well-posedness of master equations for deterministic displacement convex potential mean field games. *Communications on Pure and Applied Mathematics*, 75(12):2685–2801, 2022.

[18] Madhu Gunasingam and Ting-Kam Leonard Wong. Adapted optimal transport between gaussian processes in discrete time. *arXiv preprint arXiv:2404.06625*, 2024.

[19] Xin Guo, Xinyu Li, and Yufei Zhang. An α-potential game framework for n-player games. *arXiv preprint arXiv:2403.16962*, 2024.

[20] Yucheng Guo, Sergey Nadtochiy, and Mykhaylo Shkolnikov. Stefan problem with surface tension: uniqueness of physical solutions under radial symmetry. *arXiv preprint arXiv:2306.02969*, 2023.

[21] Camilo Hernández, Nicolás Hernández Santibáñez, Emma Hubert, and Dylan Possamaï. Closed-loop equilibria for stackelberg games: it's all about stochastic targets. *arXiv preprint arXiv:2406.19607*, 2024.

[22] Yu-Chi Ho. Optimal terminal maneuver and evasion strategy. *SIAM J. Control*, 4(3):421–428, 1966.

[23] Marc Hoffmann and Yating Liu. A statistical approach for simulating the density solution of a mckean-vlasov equation. *arXiv preprint arXiv:2305.06876*, 2023.

[24] Hui Huang, Jinniao Qiu, and Konstantin Riedl. Consensus-based optimization for saddle point problems. *SIAM Journal on Control and Optimization*, 62(2):1093–1121, 2024.

[25] Yu-Jui Huang and Zhou Zhou. Optimal equilibria for time-inconsistent stopping problems in continuous time. *Mathematical Finance*, 30(3):1103–1134, 2020.

[26] Melih Iseri and Jianfeng Zhang. Set valued hamilton-jacobi-bellman equations. *arXiv preprint arXiv:2311.05727*, 2023.

[27] Vassili N Kolokoltsov. Dynamic quantum games. *Dynamic Games and Applications*, 12(2):552–573, 2022.

[28] Vassili N Kolokoltsov. Quantum mean-field games. *The Annals of Applied Probability*, 32(3):2254–2288, 2022.

[29] Dmitry Kramkov and Mihai Sîrbu. Backward martingale transport and fitzpatrick functions in pseudo-euclidean spaces. *The Annals of Applied Probability*, 34(1B):1571–1599, 2024.

[30] Quanjun Lang, Xiong Wang, Fei Lu, and Mauro Maggioni. Interacting particle systems on networks: joint inference of the network and the interaction kernel. *arXiv preprint arXiv:2402.08412*, 2024.

[31] Zongxia Liang, Yi Xia, and Bin Zou. A two-layer stochastic game approach to reinsurance contracting and competition. *arXiv preprint arXiv:2405.06235*, 2024.

[32] Jin Ma, Gaozhan Wang, and Jianfeng Zhang. On convergence and rate of convergence of policy improvement algorithms. *arXiv preprint arXiv:2406.10959*, 2024.

[33] Alpár R Mészáros and Chenchen Mou. Mean field games systems under displacement monotonicity. *SIAM Journal on Mathematical Analysis*, 56(1):529–553, 2024.

[34] Chenchen Mou and Jianfeng Zhang. Minimal solutions of master equations for extended mean field games. *Journal de Mathématiques Pures et Appliquées*, 184:190–217, 2024.

[35] Jinniao Qiu, Antony Ware, and Yang Yang. Stochastic path-dependent volatility models for price-storage dynamics in natural gas markets and discrete-time swing option pricing. *arXiv preprint arXiv:2406.16400*, 2024.

[36] Jinniao Qiu and Jing Zhang. Stochastic differential games with random coefficients and stochastic hamilton–jacobi–bellman–isaacs equations. *The Annals of Applied Probability*, 33(2):889–930, 2023.

[37] Shanjian Tang and Jianjun Zhou. Viscosity solutions of second order path-dependent partial differential equations and applications. *arXiv preprint arXiv:2405.06309*, 2024.

[38] Xuwei Yang and Minyi Huang. Linear quadratic mean field stackelberg games: Master equations and time consistent feedback strategies. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 171–176. IEEE, 2021.