Controlled measure-valued martingales

A viscosity solution approach

Alexander M. G. Cox Sigrid Källblad Martin Larsson Sara Svaluto-Ferro

Stochastic Mass Transport

Banff, 23rd March, 2022

University of Bath

Measure Valued Martingales

Our primary object of study will be stochastic processes taking values in the space of probability measures with an additional martingale assumption. Let $\mathcal{P}$ be the set of Probability measures on $\mathbb{R}$, then:

Definition

A (Probability) Measure Valued Martingale (MVM) is a $\mathcal{P}$-valued stochastic process $\xi = (\xi_t)_{t \geq 0}$ such that $\xi(\varphi)$ is a real-valued martingale for every $\varphi \in C_b$.

Canonical example of MVM: Let X_T be an integrable $\mathbb{R}$-valued $\mathcal{F}_T$-measurable r.v., on some filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$. Then:

$$\xi_t(A) := \mathbb{P}(X_T \in A | \mathcal{F}_t)$$

is an MVM.

Aim of the Talk

We want to consider stochastic control problems of the following form:

$$\text{minimize } E\left[\int_0^\infty e^{-\beta t} c(\xi_t)dt\right]$$

over (some specified subset of) MVMs ξ_t with initial value $\xi_0 = \mu$. Here c is some cost function.

In fact, will typically restrict the class of MVMs via a control ρ, so that the evolution of ξ is determined by the control ρ. Then we will want to understand the value function:

$$v(\mu) := \inf\left\{E\left[\int_0^\infty e^{-\beta t} c(\xi_t, \rho_t)dt\right] : (\xi, \rho) \text{ admissible control}, \xi_0 = \mu\right\}$$

Aim: Make sense of this equation!

Aim of the Talk

We want to consider stochastic control problems of the following form:

$$\text{minimize } E\left[\int_0^\infty e^{-\beta t} c(\xi_t)dt\right]$$

over (some specified subset of) MVMs ξ_t with initial value $\xi_0 = \mu$. Here c is some cost function.

In fact, will typically restrict the class of MVMs via a control ρ, so that the evolution of ξ is determined by the control ρ. Then we will want to understand the value function:

$$v(\mu) := \inf\left\{E\left[\int_0^\infty e^{-\beta t} c(\xi_t, \rho_t)dt\right] : (\xi, \rho) \text{ admissible control, } \xi_0 = \mu\right\}$$

Aim: Make sense of this equation!

Aim of the Talk

We want to consider stochastic control problems of the following form:

$$\text{minimize } E\left[\int_0^\infty e^{-\beta t} c(\xi_t) dt\right]$$

over (some specified subset of) MVMs ξ_t with initial value $\xi_0 = \mu$. Here c is some cost function.

In fact, will typically restrict the class of MVMs via a control ρ, so that the evolution of ξ is determined by the control ρ. Then we will want to understand the value function:

$$v(\mu) := \inf\left\{E\left[\int_0^\infty e^{-\beta t} c(\xi_t, \rho_t) dt\right] : (\xi, \rho) \text{ admissible control, } \xi_0 = \mu\right\}$$

Aim: Make sense of this equation!

Examples: 1(a). 'Martingale Optimal Transport'

Find martingale M with $M_T \sim \mu$, $M_0 = \int x \, \mu(\mathrm{d}x)$, to max/minimise path functional of the process, e.g. the average:

$$\mathbb{E}\left[F\left(\frac{1}{T}\int_0^T M_s \, \mathrm{d}s\right)\right] = \mathbb{E}\left[F\left(\frac{1}{T}\int_0^T \int x \, \xi_s(\mathrm{d}x) \, \mathrm{d}s\right)\right]$$

where $\xi_t(A) = \mathbb{P}(M_T \in A | \mathcal{F}_t)$.

- Assume here a trivial initial law.
- Conditioning translates a terminal condition into an initial condition.

Examples: 1(b). Model-independent Option Pricing

Common financial problem: given the current prices of vanilla call options, and with minimal assumptions about dynamics of the underlying asset.

Canonical approach: Discounted asset price S is a martingale under risk-neutral measure, call prices $\implies$ law of S_T, say μ. Optimise over risk-neutral models with S a martingale and $S_T \sim \mu$ to get model independent bounds. $\rightsquigarrow$ Martingale Optimal Transport.

E.g. Asian Option pays holder a function of the average of an asset's value between time 0 and time T: i.e. holder receives $F(A_T)$, where $A_T = \frac{1}{T}\int_0^T S_r \,\mathrm{d}r$.

Equivalent to previous formulation!

Examples: 1(b). Model-independent Option Pricing

Common financial problem: given the current prices of vanilla call options, and with minimal assumptions about dynamics of the underlying asset.

Canonical approach: Discounted asset price S is a martingale under risk-neutral measure, call prices $\implies$ law of S_T, say μ. Optimise over risk-neutral models with S a martingale and $S_T \sim \mu$ to get model independent bounds. $\sim$ Martingale Optimal Transport.

E.g. Asian Option pays holder a function of the average of an asset's value between time 0 and time T: i.e. holder receives $F(A_T)$, where $A_T = \frac{1}{T}\int_0^T S_r \, \mathrm{d}r$.

Equivalent to previous formulation!

Examples: 1(b). Model-independent Option Pricing

- In C.-Källblad '17, it was shown that this problem can be solved dynamically by treating the terminal condition ($S_T \sim \mu$) as a state variable.
- The (risk-neutral) martingale condition on the call prices is exactly the constraint that the 'state variable', which is the conditional value of the law of S_T at time t, is an MVM.
- Method generalises (with some simple modifications) to large class of options. (e.g. Bayraktar, C., Stoev '18).
- Option payoffs may depend on future call prices (e.g. VIX-based options), or may want to constrain dynamics of the call prices.

Examples: 2. Optimal Skorokhod Embedding Problem

Given a measure μ, and Brownian motion B, the Skorokhod Embedding problem (SEP) is to find a stopping time τ such that $B_\tau \sim \mu$, $(B_{t\wedge\tau})_{t\geq 0}$ is u.i..

The Optimal SEP is to maximise some path functional F over all solutions to the SEP:

$$\text{maximise } \mathbb{E}\left[F\left((B_s)_{s\leq\tau}\right)\right]$$

over stopping times τ solving the SEP. A common sub-class of problems is when F is invariant to time-change.

In C.-Källblad, showed there is a one-to-one correspondence between the set of MVMs starting at μ which terminate, that is, $\xi_s \to \xi_\infty \in \mathcal{P}^s := \{\mu : \mu = \delta_x, \text{ some } x \in \mathbb{R}\}$ and set of uniformly integrable martingales with terminal law μ. Can map such martingales to stopped Brownian motion via a time-change argument $\to$ equivalence (up to time-change) between MVMs and solutions to SEP.

Examples: 2. Optimal Skorokhod Embedding Problem

Given a measure μ, and Brownian motion B, the Skorokhod Embedding problem (SEP) is to find a stopping time τ such that $B_\tau \sim \mu$, $(B_{t\wedge\tau})_{t\geq 0}$ is u.i..

The Optimal SEP is to maximise some path functional F over all solutions to the SEP:

$$\text{maximise } \mathbb{E}\left[F\left((B_s)_{s\leq\tau}\right)\right]$$

over stopping times τ solving the SEP. A common sub-class of problems is when F is invariant to time-change.

In C.-Källblad, showed there is a one-to-one correspondence between the set of MVMs starting at μ which terminate, that is, $\xi_s \to \xi_\infty \in \mathcal{P}^s := \{\mu : \mu = \delta_x, \text{ some } x \in \mathbb{R}\}$ and set of uniformly integrable martingales with terminal law μ. Can map such martingales to stopped Brownian motion via a time-change argument $\to$ equivalence (up to time-change) between MVMs and solutions to SEP.

Examples: 3. Zero Sum Games with Incomplete Information

- Setup due to Cardaliaguet and Rainer ('09, '12, . . .), based on earlier work of Aumann and Maschler ('95).
- Two player, zero sum game. Reward of game depends on a parameter θ which is known to player I, unknown to player II.
- Player II has a prior belief of the parameter θ, and learns about the parameter through the actions of player I.
- At time t, player II will update her belief about law of θ to ξ_t, where ξ is then an MVM.
- Player I chooses their actions to control ξ_t to maximise their payoff from the game.

Examples: 4. Bayesian Search Problem

- Imagine a Poisson process N on $\mathbb{R}_+ \times \mathbb{R}$ with intensity $\mathrm{d}t \times (\mathrm{Leb}(\mathrm{d}x) + \alpha\delta_y(\mathrm{d}x))$, where y is an unknown location we wish to find, with prior distribution μ.
- At time t we can centre our search on the location y_t, and we will observe a counting process which counts each point of N with probability $\gamma(z - y_t)$, where γ is a symmetric process which decreases away from zero.
- We scale the problem, increasing the rate of the Poisson process, and scaling the signal-noise ratio α to get a meaningful limit. Expect Brownian scaling in the limit.
- Our belief in the location of the true value y will be a controlled, measure-valued process, ξ_t, and in fact, an MVM, with $\xi_0 = \mu$.

Examples: 4. Bayesian Search Problem

- Search is stopped at a random exponential time. Minimise the variance of ξ at stopping:

$$\text{minimise} \int_0^\infty e^{-\kappa t} \mathbb{V}\mathrm{ar}(\xi_t) \, \mathrm{d}t$$

  where the minimisation takes place over the class of controls, (y_t).

- In the Brownian scaling, we can calculate:

$$\mathrm{d}\xi_t(f) = \alpha\Big(\int f(y)\gamma(y - y_t)\xi_t(\mathrm{d}y) - \int f(y)\xi_t(\mathrm{d}y) \int \gamma(y - y_t)\xi_t(\mathrm{d}y)\Big)\mathrm{d}W_t$$

  for the (controlled) dynamics of the posterior measure.

Existing Literature

- Lots of recent work on stochastic control of McKean-Vlasov equations, but note that the dynamics of our measures are quite different (no spatial motion, for example). E.g. Cosso et al '21; Burzoni et al '21; Talbi et al '21.
- Similarly, (mostly) old literature on controlled filtering equations. (E.g. Gozzi, Świeck '00), but these seem to rely on embedding the problem into a 'nice' function space via densities. Our approach preserves the probability measure of the original state. Recent related work: Martini, '21,'22.
- Study of measure-valued processes has a long history, e.g. martingale measures were introduced by Dawson '93.
- Eldan '16 introduced a measure-based construction of solutions to the SEP. Connections to Stochastic Localisation?

Properties of MVMs

MVMs have some nice properties:

- Let $\mathcal{P}_p = \{\mu \in \mathcal{P} : \int |x|^p \, \mu(\mathrm{d}x) < \infty\}$, equipped with Wasserstein p-metric. If ξ is an MVM with $\xi_0 \in \mathcal{P}_p$, then $\xi_t \in \mathcal{P}_p$. Moreover, if ξ has weakly continuous trajectories, then the trajectories are continuous in $\mathcal{P}_p$.
- Support of MVMs are decreasing:

  $$t \geq s \implies \mathrm{supp}(\xi_t) \subseteq \mathrm{supp}(\xi_s).$$

  But not case that $t \geq s \implies \xi_t \ll \xi_s$!
- More generally, if $\xi_0 \in \mathcal{P}_2$, then the variance is a supermartingale:

  $$\mathbb{V}\mathrm{ar}(\xi_t) = \int x^2 \, \xi_t(\mathrm{d}x) - (\mathbb{M}(\xi_t))^2$$

  where we write $\mathbb{M}(\mu) = \int x \, \mu(\mathrm{d}x)$.
- Continuous MVMs can be localised in compact sets! [By De La Vallée-Poussin]

Outline of results

Main results follow ‘classical’ structure:

1. Define an appropriate class of controlled MVMs: what does ‘control’ mean?
2. Prove the Dynamic Programming Principle for this class of controls.
3. Prove an Itô formula for MVMs: characterise martingales, setup HJB equation. Verification for ‘smooth’ value functions.
4. Introduce appropriate notion of viscosity solution: show value function satisfies HJB in an appropriate weak sense.
5. Prove comparison theorem for viscosity solutions: show there is a unique viscosity solution for HJB (which is then the value function).

Specifying the Dynamics of the MVM

Suppose ξ is an MVM on a space whose filtration is generated by a Brownian motion W. For any $\varphi \in C_b$, the martingale representation theorem yields

$$\xi_t(\varphi) = \xi_0(\varphi) + \int_0^t \sigma_s(\varphi) dW_s \qquad (1)$$

for some p. m. process $\sigma(\varphi)$ with $\int_0^t \sigma_s(\varphi)^2 ds < \infty$ for all t.

Can construct σ so that $\sigma_t(\varphi) = \int \varphi(x)\sigma_t(\mathrm{d}x)$, and typically might have $\sigma_t(\mathrm{d}x) \ll \xi_t(\mathrm{d}x)$, and $\sigma_t(1) = 0$ since $\xi_t(1) = 1$.

This implies (Yor '85, '12) existence of a function ρ such that

$$\sigma_t(\varphi) = \xi_t(\varphi\rho_t) - \xi_t(\varphi)\xi_t(\rho_t) \text{ for all } \varphi \in C_b.$$

With the notation $\mathbb{C}\mathrm{ov}_\mu(\varphi, \psi) = \mu(\varphi\psi) - \mu(\varphi)\mu(\psi)$, (1) becomes

$$\boxed{\xi_t(\varphi) = \xi_0(\varphi) + \int_0^t \mathbb{C}\mathrm{ov}_{\xi_s}(\varphi, \rho_s) dW_s.} \qquad \text{(MVM-SDE)}$$

We take the process ρ to be the control of the MVM.

Specifying the Dynamics of the MVM

Suppose ξ is an MVM on a space whose filtration is generated by a Brownian motion W. For any $\varphi \in C_b$, the martingale representation theorem yields

$$\xi_t(\varphi) = \xi_0(\varphi) + \int_0^t \sigma_s(\varphi) dW_s \tag{1}$$

for some p. m. process $\sigma(\varphi)$ with $\int_0^t \sigma_s(\varphi)^2 ds < \infty$ for all t.

Can construct σ so that $\sigma_t(\varphi) = \int \varphi(x)\sigma_t(\mathrm{d}x)$, and typically might have $\sigma_t(\mathrm{d}x) \ll \xi_t(\mathrm{d}x)$, and $\sigma_t(1) = 0$ since $\xi_t(1) = 1$.

This implies (Yor '85, '12) existence of a function ρ such that

$$\sigma_t(\varphi) = \xi_t(\varphi\rho_t) - \xi_t(\varphi)\xi_t(\rho_t) \text{ for all } \varphi \in C_b.$$

With the notation $\mathbb{C}\mathrm{ov}_\mu(\varphi, \psi) = \mu(\varphi\psi) - \mu(\varphi)\mu(\psi)$, (1) becomes

$$\boxed{\xi_t(\varphi) = \xi_0(\varphi) + \int_0^t \mathbb{C}\mathrm{ov}_{\xi_s}(\varphi, \rho_s) dW_s.} \tag{MVM-SDE}$$

We take the process ρ to be the control of the MVM.

Specifying the Dynamics of the MVM

Suppose ξ is an MVM on a space whose filtration is generated by a Brownian motion W. For any $\varphi \in C_b$, the martingale representation theorem yields

$$\xi_t(\varphi) = \xi_0(\varphi) + \int_0^t \sigma_s(\varphi) dW_s \tag{1}$$

for some p. m. process $\sigma(\varphi)$ with $\int_0^t \sigma_s(\varphi)^2 ds < \infty$ for all t.

Can construct σ so that $\sigma_t(\varphi) = \int \varphi(x)\sigma_t(\mathrm{d}x)$, and typically might have $\sigma_t(\mathrm{d}x) \ll \xi_t(\mathrm{d}x)$, and $\sigma_t(1) = 0$ since $\xi_t(1) = 1$.

This implies (Yor '85, '12) existence of a function ρ such that

$$\sigma_t(\varphi) = \xi_t(\varphi\rho_t) - \xi_t(\varphi)\xi_t(\rho_t) \text{ for all } \varphi \in C_b.$$

With the notation $\mathbb{C}\mathrm{ov}_\mu(\varphi, \psi) = \mu(\varphi\psi) - \mu(\varphi)\mu(\psi)$, (1) becomes

$$\boxed{\xi_t(\varphi) = \xi_0(\varphi) + \int_0^t \mathbb{C}\mathrm{ov}_{\xi_s}(\varphi, \rho_s) dW_s.} \tag{MVM-SDE}$$

We take the process ρ to be the control of the MVM.

Weak solutions to MVM-SDE

We now need to discuss what we mean by a solution of the control problem. We work with weak formulations:

Definition

A weak solution of (MVM-SDE) is a tuple $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P}, W, \xi, \rho)$, where $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, P)$ is a filtered probability space, W is a standard Brownian motion on this space, ξ is a continuous MVM, and ρ is a progressively measurable function on $\Omega \times \mathbb{R}_+ \times \mathbb{R}$ such that for every $\varphi \in C_b$, $P \otimes dt$-a.e.,

$$\xi_t(|\rho_t|) < \infty, \quad \int_0^t \mathbb{C}\mathrm{ov}_{\xi_s}(\varphi, \rho_s)^2 ds < \infty,$$

and (MVM-SDE) holds, that is,

$$\xi_t(\varphi) = \xi_0(\varphi) + \int_0^t \mathbb{C}\mathrm{ov}_{\xi_s}(\varphi, \rho_s) dW_s.$$

Admissible Controls

We fix $p \geq 0, q \in [0, p]$ and a Polish space $\mathbb{H}$ of measurable, real functions which will contain the control. Then we say:

Definition

An *admissible control* is a weak solution (ξ, ρ) of (MVM-SDE) such that

$$\rho_t(\cdot, \omega) \in \mathbb{H}$$

and, $\mathbb{P} \otimes dt$-a.e.,

$$\int_0^t \left(\int_{\mathbb{R}} (1 + |x|^q) |\rho_s(x) - \xi_s(\rho_s)| \xi_s(dx) \right)^2 ds < \infty.$$

NB: Can guarantee second conditions by placing growth bounds on $\mathbb{H}$.

Can also have suitable state-dependent restriction on the control.

Control Problem and DPP

We consider the following control problem. In addition to the action space $\mathbb{H}$, fix a measurable cost function

$$c\colon \mathcal{P}_p \times \mathbb{H} \to \mathbb{R} \cup \{+\infty\}$$

and a discount rate $\beta \geq 0$. For $\mu \in \mathcal{P}_p$ the value function is given by

$$v(\mu) = \inf \left\{ E\left[\int_0^\infty e^{-\beta t} c(\xi_t, \rho_t) dt\right] : (\xi, \rho) \text{ admissible control, } \xi_0 = \mu \right\}.$$

Theorem (Dynamic Programming Principle)

Let τ be a bounded stopping time on $C(\mathbb{R}_+, \mathcal{P}_p)$. For any $\mu \in \mathcal{P}_p$, the value function v satisfies

$$v(\mu) = \inf_{(\xi,\rho)} \mathbb{E}\left[e^{-\beta\tau(\xi)} v(\xi_{\tau(\xi)}) + \int_0^{\tau(\xi)} e^{-\beta t} c(\xi_t, \rho_t) dt \right],$$

where the infimum extends over all admissible controls (ξ, ρ) with $\xi_0 = \mu$.

Proof: Using general framework of Žitković '14.

Existence of non-trivial solutions

One important question is: given a choice of the control ρ, does (MVM-SDE) have a solution? This is non-trivial, but the answer is yes!

Theorem (Global Existence of solutions)

For any measurable function $\bar{\rho}\colon \mathbb{R} \to \mathbb{R}$ and any $\mu \in \mathcal{P}$, there exists a weak solution (ξ, ρ) of (MVM-SDE) *such that $\xi_0 = \mu$ and $\rho_t = \bar{\rho}$ for all t.*

Proof via a careful construction argument.

Towards an Itô formula: derivatives of functions of measures

Next want to consider measure-valued functions, and their derivatives:

Definition (C.f. Carmona and Delarue '18)

Let $p \geq 0$. A function $f : \mathcal{P}_p \to \mathbb{R}$ is said to belong to $C^1(\mathcal{P}_p)$ if there is a continuous function $(x, \mu) \mapsto \frac{\partial f}{\partial \mu}(x, \mu)$ from $\mathbb{R} \times \mathcal{P}_p$ to $\mathbb{R}$, called (a version of) the derivative of f, with the following properties.

- locally uniform p-growth: for every compact set $K \subset \mathcal{P}_p$, there is a constant c_K such that for all $x \in \mathbb{R}$ and $\mu \in K$,

$$\left| \frac{\partial f}{\partial \mu}(x, \mu) \right| \leq c_K(1 + |x|^p),$$

- fundamental theorem of calculus: for every $\mu, \nu \in \mathcal{P}_p$,

$$f(\nu) - f(\mu) = \int_0^1 \int_{\mathbb{R}} \frac{\partial f}{\partial \mu}(x, t\nu + (1-t)\mu)(\nu - \mu)(dx)dt.$$

Similar definition for C^2, etc.

Itô formula for MVMs

Theorem (Itô's Formula)

Let (ξ, ρ) *be a weak solution of* (MVM-SDE)*, where* ξ *takes values in* $\mathcal{P}_p$ *for some fixed* $p \geq 0$*. Let* $q \in [0, p]$ *and assume that,* $P \otimes dt$*-a.e.,*

$$\int_0^t \left(\int_{\mathbb{R}} (1 + |x|^q) |\rho_s(x) - \xi_s(\rho_s)| \xi_s(dx) \right)^2 ds < \infty.$$

Then, for every f *in* $C^2(\mathcal{P}_q)$ *we have the Itô formula*

$$\boxed{\begin{aligned} f(\xi_t) = f(\xi_0) &+ \int_0^t \int_{\mathbb{R}} \frac{\partial f}{\partial \mu}(x, \xi_s) \sigma_s(dx) dW_s \\ &+ \frac{1}{2} \int_0^t \int_{\mathbb{R} \times \mathbb{R}} \frac{\partial^2 f}{\partial \mu^2}(x, y, \xi_s) \sigma_s(dx) \sigma_s(dy) ds, \end{aligned}}$$

where we write $\sigma_s(dx) = (\rho_s(x) - \xi_s(\rho_s)) \xi_s(dx)$*.*

Proof: See Sigrid's talk.

HJB Formulation

We expect v to be a solution (in some sense) of the following HJB equation:

$$\beta u(\mu) + \sup_{\rho \in \mathbb{H}} \{-c(\mu,\rho) - Lu(\mu,\rho)\} = 0, \qquad \mu \in \mathcal{P}_p \setminus \mathcal{P}^s$$

$$u(\mu) = c(x)/\beta, \qquad \mu = \delta_x \in \mathcal{P}^s$$

where

$$c(x) = \inf_{\rho \in \mathbb{H}} c(\delta_x, \rho)$$

and the operator L is given by

$$Lf(\mu,\rho) = \frac{1}{2} \int_{\mathbb{R}\times\mathbb{R}} \frac{\partial^2 f}{\partial \mu^2}(x,y,\mu)\sigma(dx)\sigma(dy)$$

with $\sigma(dx) = (\rho(x) - \mu(\rho))\mu(dx)$.

The boundary condition can be understood as follows: since an MVM starting at a Dirac measure δ_x stays there, v must satisfy $v(\delta_x) = c(x)/\beta$.

Main Result I: Viscosity Characterisation

Theorem

Fix a set of actions $\mathbb{H}$*, a discount rate* $\beta > 0$*, and a cost function* $c : \mathcal{P}_p \times \mathbb{H} \to \mathbb{R} \cup \{\infty\}$*. Assume that*

1. *there is a constant* $\kappa \in (0, \infty)$ *such that* $|\rho(x)| \leq \kappa(1 + |x|^p)$ *holds for all* $x \in \mathbb{R}$ *and* $\rho \in \mathbb{H} \cap C_c(\mathbb{R})$*;*
2. $\mu \mapsto c(\mu, \rho)$ *is upper semi-continuous for every* $\rho \in \mathbb{H} \cap C_c(\mathbb{R})$*;*
3. *for every* $\mu \in \mathcal{P}_p$ *and every* $f \in C^2(\mathcal{P}_q)$*,*

$$\sup_{\rho \in \mathbb{H}} \{-c(\mu, \rho) - Lf(\mu, \rho)\} = \sup_{\rho \in \mathbb{H} \cap C_c(\mathbb{R})} \{-c(\mu, \rho) - Lf(\mu, \rho)\};$$

Then the value function $v\colon \mathcal{P}_p \to \mathbb{R}$ *is a viscosity solution of the HJB equation.*

Viscosity Solutions I: Restricting limits

Since MVMs have decreasing support, define a partial order $\preceq$ on $\mathcal{P}_p$ by

$$\mu \preceq \nu \quad \iff \quad \mathrm{supp}(\mu) \subset \mathrm{supp}(\nu).$$

Note that MVMs are decreasing with respect to this order. So effective state space for an MVM starting at a measure $\bar{\mu} \in \mathcal{P}_p$ is the set

$$D_{\bar{\mu}} = \{\mu \in \mathcal{P}_p \colon \mu \preceq \bar{\mu}\}.$$

In particular, for any $u \colon \mathcal{P}_p \to \overline{\mathbb{R}}$, the restriction of u to $D_{\bar{\mu}}$ has semicontinuous envelopes given by

$$(u|_{D_{\bar{\mu}}})^*(\mu) = \limsup_{\nu \to \mu,\, \nu \preceq \bar{\mu}} u(\nu)$$
$$(u|_{D_{\bar{\mu}}})_*(\mu) = \liminf_{\nu \to \mu,\, \nu \preceq \bar{\mu}} u(\nu)$$

for all $\mu \preceq \bar{\mu}$.

Viscosity Solutions I: Definition

For any test function $f \in C^2(\mathcal{P}_q)$, define $H(\cdot; f): \mathcal{P}_p \to \overline{\mathbb{R}}$ by

$$H(\mu; f) = \beta f(\mu) + \sup_{\rho \in \mathbb{H}} \left\{ -c(\mu, \rho) - Lf(\mu, \rho) \right\}.$$

We can now state the definition of viscosity solution.

Definition

Consider a function $u: \mathcal{P}_p \to \overline{\mathbb{R}}$.

- u is a viscosity subsolution if

$$\liminf_{\mu \to \bar{\mu},\, \mu \preceq \bar{\mu}} H(\mu; f) \le 0$$

  holds for all $\bar{\mu} \in \mathcal{P}_p \setminus \mathcal{P}^s$ and $f \in C^2(\mathcal{P}_q)$ such that $f(\bar{\mu}) = \limsup_{\mu \to \bar{\mu},\, \mu \preceq \bar{\mu}} u(\mu)$ and $f(\mu) \ge u(\mu)$ for all $\mu \preceq \bar{\mu}$.

Viscosity Solutions I: Definition

For any test function $f \in C^2(\mathcal{P}_q)$, define $H(\cdot; f)\colon \mathcal{P}_p \to \overline{\mathbb{R}}$ by

$$H(\mu; f) = \beta f(\mu) + \sup_{\rho \in \mathbb{H}} \{-c(\mu, \rho) - Lf(\mu, \rho)\}.$$

We can now state the definition of viscosity solution.

Definition

Consider a function $u\colon \mathcal{P}_p \to \overline{\mathbb{R}}$.

- u is a viscosity supersolution if

$$\limsup_{\mu \to \bar{\mu},\, \mu \preceq \bar{\mu}} H(\mu; f) \geq 0$$

holds for all $\bar{\mu} \in \mathcal{P}_p \setminus \mathcal{P}^s$ and $f \in C^2(\mathcal{P}_q)$ such that $f(\bar{\mu}) = \liminf_{\mu \to \bar{\mu},\, \mu \preceq \bar{\mu}} u(\mu)$ and $f(\mu) \leq u(\mu)$ for all $\mu \preceq \bar{\mu}$.

Viscosity Solutions I: Definition

For any test function $f \in C^2(\mathcal{P}_q)$, define $H(\cdot; f)\colon \mathcal{P}_p \to \overline{\mathbb{R}}$ by

$$H(\mu; f) = \beta f(\mu) + \sup_{\rho \in \mathbb{H}} \{-c(\mu, \rho) - Lf(\mu, \rho)\}.$$

We can now state the definition of viscosity solution.

Definition

Consider a function $u\colon \mathcal{P}_p \to \overline{\mathbb{R}}$.

- u is a viscosity solution if it is both a viscosity subsolution and a viscosity supersolution.

Lemma: Every classical solution is a viscosity solution.

Main Result II: Comparison

Need to show that viscosity solutions are unique:

Theorem

Let $\beta > 0$, and suppose that the cost function c and the action space $\mathbb{H}$ satisfy the following conditions:

1. *$\mu \mapsto c(\mu, \rho)$ is continuous on $\mathcal{P}(\{x_1, ..., x_N\})$ uniformly in $\rho \in \mathbb{H}$ for any $N \in \mathbb{N}$ and $x_1, ..., x_N \in \mathbb{R}$;*
2. *the set $\{\rho(x) - \rho(0) \colon \rho \in \mathbb{H}\}$ is bounded for every $x \in \mathbb{R}$.*

Let $u, v \in C(\mathcal{P}_p)$ be a viscosity sub- and supersolution, respectively. If $u \leq v$ on $\mathcal{P}^s$, then $u \leq v$ on $\mathcal{P}_p$.

Value function is uniquely characterised as the viscosity solution of the HJB equation!

Remarks

- Can incorporate state dependent constraints by modifying the cost. Let $A \subseteq \mathcal{P}_p \times \mathbb{H}$ be the state constraint, so we require $(\xi_t, \rho_t) \in A$, and suppose A is open. Then we can set the cost function to $+\infty$ on $A^{\complement}$ to recover a constrained problem.
- Problem generalises to MVMs on $\mathbb{R}^d$.
- Open questions: how general is the requirement that the MVM solves (MVM-SDE)?

Example: Explicitly solvable control problems

Example

Fix $p \geq 4$, $q = 1$, a state dependent set of actions

$$\mathbb{H}(\mu) := \{\rho \in \mathbb{H} \colon \mathbb{V}\mathrm{ar}_\mu(\rho) \leq \mathbb{V}\mathrm{ar}(\mu)\}$$

for some $\mathbb{H}$ such that $\mathrm{id} \in \mathbb{H}$, and a discount rate $\beta > 0$. Define

$$c(\mu, \rho) := 2\mathbb{V}\mathrm{ar}(\mu)^2 - \beta\mathbb{M}(\mu)^2.$$

Then the value function is the unique continuous viscosity solution of the HJB equation and is given by

$$\mathbb{M}(\mu)^2 = \inf\left\{E\left[\int_0^\infty e^{-\beta t} c(\xi_t, \rho_t) dt\right] : (\xi, \rho) \text{ admissible control, } \xi_0 = \mu\right\}.$$

Moreover, there exists an optimal control (ξ^*, ρ^*) satisfying $\xi_s^*(\rho_s^*) = \mathbb{M}(\xi_s^*)$ for a.e. $s \geq 0$ (e.g. $\rho_s^*(x) = x$).

This solution already appeared in the context of the SEP in Eldan '16.

Conclusions

- Consider Stochastic Control problems in the space of MVMs: can describe a wide range of interesting control problems.
- Develop stochastic representation for the controlled process, with corresponding Itô formula.
- Construct appropriate notion of viscosity solution, show value function is unique viscosity solution to the HJB equation.
- Can derive optimal behaviour in some simple control problems.