Physics Letters A 384 (2020) 126608

Contents lists available at ScienceDirect

Physics Letters A

www.elsevier.com/locate/pla

Universal behavior in non-stationary Mean Field Games

Thibault Bonnemain [a,b], Thierry Gobron [b,c,*], Denis Ullmo [a]

[a] Université Paris-Saclay, CNRS, LPTMS, 91405 Orsay, France
[b] LPTM, CNRS, Université Cergy-Pontoise, 95302 Cergy-Pontoise, France
[c] Laboratoire Paul Painlevé, Université de Lille, 59655 Villeneuve d'Ascq, France

ARTICLE INFO

Article history:
Received 23 January 2020
Received in revised form 17 May 2020
Accepted 22 May 2020
Available online 27 May 2020
Communicated by M. Perc

Keywords:
Mean Field Games
Econophysics
Integrable systems
Electrostatics
Poisson and Laplace equations
Boundary-value problems

ABSTRACT

Mean Field Games provide a powerful framework to analyze the dynamics of a large number of controlled objects in interaction. Though these models are much simpler than the underlying differential games they describe in some limit, their behavior is still far from being fully understood. When the system is confined, a notion of "ergodic state" has been introduced that characterizes most of the dynamics for long optimization times. Here we consider a class of models without such an ergodic state, and show the existence of a scaling solution that plays a similar role. Its universality and scaling behavior can be inferred from a mapping to an electrostatic problem.

Mean Field Games are a powerful framework introduced about a decade ago by Lasry and Lions [1] as an alternative approach to differential game theory when the number of agents becomes large (see [2] for a suitable introduction to the field for physicists). Their applications are numerous, ranging from finance [3,4] and economy [5,6] to engineering sciences [7,8], and wherever one has to deal with optimization issues for many coupled subsystems. On a quite general basis, a Mean Field Game involves a set of N players (or agents) which are characterized by a continuous "state variable" $\mathbf{X}^i \in \mathbb{R}^d$, $i = 1 \dots N$, which, depending on the context, may represent a physical position, the amount of resources owned by a company, the house temperature in a network of controlled heaters, etc. These state variables evolve on a time interval $[0, T]$ according to some controlled dynamics, which we assume here to be described by a linear, d-dimensional, Langevin equation, $d\mathbf{X}_t^i = \mathbf{a}_t^i dt + \sigma d\mathbf{W}_t^i$, where each component of $\mathbf{W}^i$ is an independent white noise of variance one, σ is a constant, and the "control parameter" is the velocity $\mathbf{a}_t^i$. This control is adjusted in time by the agent i in order to minimize a cost functional over the remaining time to play, which in the simplest case can be assumed of the form

$$c[\mathbf{a}^{\mathbf{i}}](\mathbf{x}_t^i, t) = \mathbb{E}\left[\int_t^T \left(\frac{\mu}{2}(\mathbf{a}_t^i)^2 - V[m_\tau](\mathbf{X}_\tau^i)\right) d\tau + c_T(\mathbf{X}_T^i)\middle| \mathbf{X}_t^i = \mathbf{x}_t^i\right]. \quad (1)$$

In (1), $c_T(\mathbf{x})$ is the *terminal cost*, an additional cost player i would pay if his state at the end of the game, $\mathbf{X}_i(T)$, is equal to $\mathbf{x}$, μ is a positive constant, akin to a mass in physics, and $V[m_t](\mathbf{x})$ is a functional of the empirical density at time t $m_t(\mathbf{x}) = \frac{1}{N}\sum_j \delta(\mathbf{x} - \mathbf{X}^j(t))$, through which the agents' optimization problems are coupled. We shall assume $V[m_t](\mathbf{x})$ takes the simple form

$$V[m_t](\mathbf{x}) = U_0(\mathbf{x}) + g\, m_t(\mathbf{x})\,. \quad (2)$$

For a very large number of players, like in a mean field theory, the fluctuations of the empirical density are neglected and $m_t(\mathbf{x})$ becomes a deterministic quantity governed by a Fokker Planck equation. Furthermore, the optimization problems decouple and the optimal value of the cost (1) (the "value function") for the agent i becomes a function of the variable $\mathbf{x}_t^i$, solution of an Hamilton-Jacobi Bellman equation [9]. The resulting model is called a Mean Field Game and can be defined as a pair of coupled equations, describing the (forward in time) evolution of the density of players $m(\mathbf{x}, t)$ and the (time backward) evolution of the value function $u(\mathbf{x}, t)$. In the simple case we consider here the system of deterministic equations takes the form [2,10]

* Corresponding author.
E-mail address: thierry.gobron@m4x.org (T. Gobron).

https://doi.org/10.1016/j.physleta.2020.126608
0375-9601/

$$\begin{cases} \partial_t m - \dfrac{1}{\mu}\nabla\,[m\nabla u] - \dfrac{\sigma^2}{2}\Delta m = 0 & \text{(a)} \\ \partial_t u - \dfrac{1}{2\mu}\,[\nabla u]^2 + \dfrac{\sigma^2}{2}\Delta u = V[m]\,. & \text{(b)} \end{cases} \tag{3}$$

The coupling between the two PDE's comes from two parts: in the Fokker-Planck equation (3a), the optimal velocity appears in the drift term and here is proportional to the gradient of the value function, $\mathbf{a} = -\frac{1}{\mu}\nabla u$; in the Hamilton-Jacobi Bellman equation (3b) the term $V[m]$ reflects the dependence of the cost functional (1) on the density. This structure also induces rather atypical boundary conditions: the (forward) Fokker-Planck equation is associated with an initial condition $m(0,\mathbf{x}) = m_0(\mathbf{x})$ specifying the initial distribution of agents, while the terminal cost C_T in Eq. (1) imposes a final condition for the value function, $u(T,\mathbf{x}) = c_T(\mathbf{x})$. This forward-backward structure together with mixed initial-final boundary conditions leads to new challenges when trying to characterize, either analytically or numerically, solutions to this system of equations.

For a large class of settings, which includes in particular the case of repulsive interactions when the system is confined either by a stabilizing external potential or because of a geometry with fixed spatial extension, such a system exhibits an ergodic (stationary) state, independent on the boundary conditions, which can be rigorously defined in the limit $T \to \infty$ as an hyperbolic fixed point. The importance of this result, as proven by Cardialaguet et al. [11] is twofold: for finite but long enough optimization time T, the game will stay very close to this ergodic (time independent) state except possibly in its initial and final parts. Furthermore, the transient dynamics near $t = 0$ and $t = T$ completely decouple one from the other, and the mixed boundary problem simplifies accordingly: they describe how the initial and final boundary conditions match with the ergodic state and are characterized by two possibly different time scales, τ_i and τ_f, independent on T.

The problem we want to address in this letter is the behavior of such a system when an ergodic state cannot exist, for instance when the state space is unbounded and the interactions repulsive, so that any initially localized configuration will expand forever. The natural question in such conditions is whether some kind of limiting (non-stationary) regime still exists for such systems that would, at some level, play a similar role as the ergodic state in confined systems.

Considering more specifically a one-dimensional problem without "one-body" term (i.e. $U_0(x) \equiv 0$), a first indication that this is indeed the case is brought out by numerical solutions for the system (3a)–(3b) for repulsive interactions (that is $V[m_t](x) = gm(t,x)$, $g < 0$), starting from a localized initial density. Numerical results (see Fig. 1) show clearly that except for transient times near $t = 0$ and near $t = T$, the density can be accurately fitted on an inverted parabolic shape

$$m(x,t) = \begin{cases} \dfrac{1}{N(t)}(z(t)^2 - x^2) & \text{if } |x| \le z(t) \\ 0 & \text{otherwise}\,, \end{cases} \tag{4}$$

where the prefactor is found to be $N(t) = \frac{4}{3}z(t)^3$, as total mass conservation requires. Assuming this shape as an ansatz, the evolution reduces to that of the scaling factor $z(t)$, which is for most times found to grow as a power law (see left insert in Fig. 1)

$$z(t) \sim t^{2/3}\,. \tag{5}$$

These features appear for sufficiently long optimization time T, and are essentially independent of the initial and final boundary conditions provided the initial distribution has a bounded extension and the final cost is close to zero everywhere.

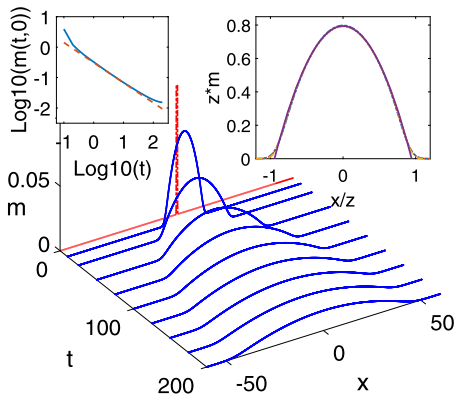

Fig. 1. Numerical simulation of the time evolution described by the system (3a)–(3b) in one space dimension with $T = 200$, $g = -2$, $\sigma = 0.5$ and $\mu = 1$. The initial distribution is a Gaussian of variance 0.1 centered at the origin, and the terminal cost is $c_T(x) = 0$. The left inset shows a log-log plot of the time evolution of the density at the origin $m(0,t)$ (full line), showing a $t^{2/3}$ scaling behavior (dashed line) at almost all times. On the right inset the rescaled density $z(t)m(x/z(t),t)$ is shown to have an invariant parabolic shape (with $z(t) = 4/(3m(0,t))$).

By analogy with the physics of Bose condensate, let us introduce a characteristic length-scale, the healing length $\nu = \mu\sigma^2/|g|$, which represent the typical distance on which interactions balances quantum pressure (or diffusion in our case) [12]. A necessary condition for the behavior observed in Fig. 1 to appear is the smallness of ν with respect to other length-scales. This condition will however eventually be fulfilled at one point as the extension of density distribution grows in time, possibly inducing a time shift in (5).

What those numerical results tell us is that in this setting, the notion of ergodic state has been replaced by the next best thing, namely a universal scaling solution. The goal of this paper is to understand this puzzling result, in particular the $\frac{2}{3}$ scaling exponent, and give some arguments in favor of its universality.

For this purpose, we now introduce a few formal transformations which allow to show that this result is rather natural and intuitive, once the problem is cast in its proper language. In turn, we will also gain a better understanding of this regime.

1. Hydrodynamic representation

The main idea behind the approach below comes from the deep link between Mean Field Games and the non-linear Schrödinger equation, discussed at lengths in [2], which allows for the use of several techniques developed to study Bose-Einstein condensates. One of these techniques, the Madelung substitution, is particularly well suited to deal with the small ν regime [13]. It consists in defining a velocity field $v(t,x)$

$$v(x,t) = -\frac{\nabla u}{\mu} - \frac{\sigma^2\nabla m}{2m}\,, \tag{6}$$

which maps the evolution (3a) for the density m to a simple transport/continuity equation

$$\partial_t m + \nabla(mv) = 0\,. \tag{7}$$

The evolution for the velocity field v derives from the HJB equation (3b) and reads

$$\partial_t v + \nabla\left[\frac{\sigma^4}{2\sqrt{m}}\Delta\sqrt{m} + \frac{v^2}{2} + \frac{g}{\mu}m\right] = 0\,, \tag{8}$$

and involves a $O(\sigma^4)$ term. As in the context of cold atoms, this term can be neglected as long as the characteristic length of the

system is large in front of the healing length $\nu = \mu\sigma^2/|g|$ [13], leading to what is referred to as the Thomas Fermi approximation. This weak noise limit is also one of the requirements for the appearance of the parabolic behavior described in Fig. 1. In this approximation, the $O(\sigma^4)$ term in (8) is dropped and the equations read

$$\begin{cases} \partial_t m + \nabla(mv) = 0 \\ \partial_t v + \nabla\left[\dfrac{v^2}{2} + \dfrac{g}{\mu}m\right] = 0\,. \end{cases} \tag{9}$$

These equations are formally very close to those studied for instance in the field of cold atoms [12], the main differences being the (negative) sign of g, which makes the system elliptic rather than hyperbolic, and the nature of the boundary conditions. Within this approximation, it can easily be verified that the ansatz Eq. (4) with

$$z(t) = 3\left(\frac{|g|}{4\mu}\right)^{1/3} t^{2/3}\,, \tag{10}$$

is a particular solution of the Thomas-Fermi-like equations Eq. (9). The real mystery is therefore not that such a particular solution does exist, but to understand why this $t^{2/3}$ behavior shows up at all intermediate times for a large class of boundary conditions, and in other terms, *why it is universal* in the large optimization time limit.

2. Riemann invariant and hodograph transform

To answer this question, we shall turn to an approach developed in the context of non-linear waves [14], which relies on the notions of Riemann invariants and hodograph transform. Riemann's method can be considered an extension of the method of characteristics. It amounts to finding curves (characteristics) on which some quantities (Riemann invariants) are conserved. Here, one can show that there exists a pair of Riemann invariants, $(\lambda_+(x,t), \lambda_-(x,t))$, namely $\lambda_\pm = v \pm 2i\sqrt{|g|m/\mu}$, so that (9) reads

$$\begin{cases} \partial_t\lambda_+ + \left(\dfrac{3}{4}\lambda_+ + \dfrac{1}{4}\lambda_-\right)\partial_x\lambda_+ = 0 \\ \partial_t\lambda_- + \left(\dfrac{1}{4}\lambda_+ + \dfrac{3}{4}\lambda_-\right)\partial_x\lambda_- = 0\,. \end{cases} \tag{11}$$

Though characteristics do not exist in this context (they are curves in the complex plane $\mathbb{C}^2$), this change of variables still allows us to linearize these equations using an hodograph transformation [14]. Taking the pair (λ_+, λ_-) as independent variables, we express x and t as functions of them, so that the system (11) transforms into a linear one

$$\begin{cases} \partial_- x - \beta_+\partial_- t = 0 \\ \partial_+ x - \beta_-\partial_+ t = 0\,, \end{cases} \tag{12}$$

where $\beta_\pm = \left(\frac{3}{4}\lambda_\pm + \frac{1}{4}\lambda_\mp\right)$, and where we have introduced the notation $\partial_\pm \equiv \frac{\partial}{\partial\lambda_\pm}$. This system can be readily integrated once as

$$\begin{cases} x - \beta_+ t = \omega_+ \\ x - \beta_- t = \omega_-\,, \end{cases} \tag{13}$$

with $\omega_\pm$ solution of

$$\partial_\pm\omega_\mp = -(\partial_\pm\beta_\mp)t = -\frac{1}{4}t\,. \tag{14}$$

Thus the functions $\omega_\pm$ can be expressed as derivatives of a potential $\omega_\pm = \partial_\pm\chi$, where $\chi(\lambda_+, \lambda_-)$ is solution of an Euler-Poisson-Darboux equation:

$$\partial^2_{+-}\chi - \frac{1}{2(\lambda_+ - \lambda_-)}(\partial_+\chi - \partial_-\chi) = 0\,. \tag{15}$$

The main difference with the traditional treatment of NLS that we have closely followed until now is that here the Riemann invariants (λ_+, λ_-) are complex conjugates. In terms of its real and imaginary parts, $\xi = \frac{1}{2}(\lambda_+ + \lambda_-) = v$ and $\eta = \frac{1}{2i}(\lambda_+ - \lambda_-) = 2\sqrt{|g|m/\mu}$, equation (15) becomes

$$\partial_{\xi\xi}\chi + \partial_{\eta\eta}\chi + \frac{1}{\eta}\partial_\eta\chi = 0\,, \tag{16}$$

which is the Laplace equation in cylindrical coordinates (with no angular dependence) with η and ξ as radial and axial coordinates, respectively. Equations (13) now read

$$\begin{cases} \eta t = -E_\eta \\ 2(x - \xi t) = -E_\xi\,, \end{cases} \tag{17}$$

with E_η and E_ξ the radial and axial components of the electric field, $\mathbf{E} = -\nabla\chi$. Note that even if (16) is originally a two-dimensional problem, a clear connection with electrostatics emerges when considering it a three dimensional one with axial symmetry.

3. Potential representation

Through the hodograph transform we have shown that for any potential χ, solution of the Laplace equation (16), there is a solution to the Thomas Fermi equation (9) provided that the relations (17) between x, t and the electric field $\mathbf{E}$ hold. The linear Laplace equation (and the related electrostatic problem) is clearly significantly simpler than the original non-linear hydrodynamic equations. The price to pay for that simplification is that taking into account the boundary conditions becomes highly non-trivial since the locus of the curves $t(\xi,\eta) = 0$ or $t(\xi,\eta) = T$ on which these conditions are expressed actually depend on the particular potential $\chi(\xi,\eta)$ considered.

However, since the dynamics we are interested in is associated with the spreading of the density of agents, the curves $t(\xi,\eta) = \text{const.}$ are contracting as t increases, smaller time curves including every larger time ones. If we consider χ as generated by a distribution of charge $\rho(\xi,\eta)$, Eq. (16) implies that $\rho(\xi,\eta) = 0$ between the curves $t(\xi,\eta) = 0$ and $t(\xi,\eta) = T$ but can be non-zero either near the origin (for times larger than T) or at large distance (corresponding to negative times). If the optimization time is long enough we can assume that there exists a range of times $[\tilde{t}_{\min}, \tilde{t}_{\max}]$, $0 \ll \tilde{t}_{\min}$, $\tilde{t}_{\max} \ll T$, so that for any curve $t(\xi,\eta) \in [\tilde{t}_{\min}, \tilde{t}_{\max}]$ the distributions of charges both at the origin and "at infinity" are far enough so that the effects of their moments are essentially negligible. Since the total charge at infinity contributes only by an irrelevant constant, we can give a good approximation of χ in this range of time as the potential created by a point charge Q_0 located at the origin

$$\chi(\eta,\xi) \approx \frac{Q_0}{\sqrt{\eta^2 + \xi^2}}\,, \tag{18}$$

with a relation between Q_0 and the boundary conditions of the problem yet to be determined.

The main result of this paper is that the approximation of the charge distribution by a monopole centered at the origin is precisely the observed universal behavior expressed by Eqs. (4)-(5).

In addition the conditions under which this approximation is valid provide the regime of validity of the scaling form Eqs. (4)-(5). Indeed, inserting Eq. (18) into Eq. (17) and inverting the relation between $(\xi = v, \eta = \sqrt{|g|m/\mu})$ and (x, t) readily gives

$$\begin{cases} m(t,x) = \dfrac{3\left((\mu Q_0/2g)^{2/3} z(t)^2 - x^2\right)}{4z(t)^3} \\ v(t,x) = -\dfrac{z'(t)}{z(t)} x\,, \end{cases} \tag{19}$$

for $|x| < z(t)$, with $z(t)$ given by Eq. (10). It suffices now to show that *independently of the boundary condition* we have $Q_0 = 2g/\mu$ (which is in any case required by the normalization of $m(x,t)$), to prove that Eq. (19) is in fact the numerically observed scaling form Eqs. (4)-(5).

To show this, we can simply apply Gauss's law

$$\int_{S_{\tilde{t}}} (\vec{E} \cdot \vec{n}) dS = 4\pi\, Q_0 \tag{20}$$

on a surface $S_{\tilde{t}}$ such that $t(\eta, \xi) = \tilde{t} = \text{const.}$, for any $0 \leq \tilde{t} \leq T$. Parametrizing this surface by (x, θ), with $dS = \eta\, j(x, \tilde{t}) d\theta dx$ (with $j(x,\tilde{t}) \equiv \sqrt{(\partial_x \xi)^2 + (\partial_x \eta)^2}$), and the normal $\vec{n}$ to the surface $S_{\tilde{t}}$ given by $\vec{n} = (n_\xi, n_\eta, n_\theta) = j(x,\tilde{t})^{-1}\,(\partial_x \xi, -\partial_x \eta, 0)$, we get

$$\begin{aligned} Q_0 &= \tfrac{1}{4\pi} \int_0^{2\pi} \int_{\mathbb{R}} \eta \left[2(x - \xi\tilde{t})\partial_x \eta - \eta \tilde{t} \partial_x \xi\right] d\theta dx \\ &= \tfrac{1}{2} \int_{\mathbb{R}} \left[-\tilde{t}\partial_x(\eta^2 \xi) + 2x\eta\partial_x \eta\right] dx\,. \end{aligned} \tag{21}$$

If we assume η to decrease sufficiently fast with x, the first, time dependent, term integrates to zero. This was to be expected because no matter the time $\tilde{t}$, as long as $0 < \tilde{t} < T$, the total charge included in $S_{\tilde{t}}$ is the same: Q_0 is by construction *a constant of the motion*. Integrating by part and recalling that $\eta = 2\sqrt{|g|m/\mu}$, Eq. (21) yields $Q_0 = \frac{2g}{\mu} \int_{\mathbb{R}} m dx$, which, because of the normalization condition on $m(x,t)$, is the required result.

4. Conclusion

Introducing a potential representation for the [1d potential-free] Mean Field Game problem Eqs. (3a)–(3b), we have shown that the remarkable, and a priori quite puzzling *universal scaling form* we have observed numerically for large optimization times can be derived in a very natural way. In this representation, it appears as the simple fact that the related potential χ, sufficiently far away from the charge distribution which creates it, can be well approximated by the potential generated by a single monopole. The condition of being sufficiently far from both the charges near the origin and the one near infinity is reflected in the MFG problem by considering times far from both $t = 0$ and $t = T$, which is of course only possible in the limit of very long optimization time. In some sense, we have thus been allowed to replace the notion of ergodic state by what could be thus called "an ergodic scaling form" in a situation where a genuine ergodic state cannot exist.

At a more general level, the potential representation of this Mean Field Game underlies the *integrability* of the hydrodynamical equations (9). Beyond the simple monopolar approximation Eq. (18), we can construct a complete multipolar expansion for the potential χ, and it can be seen easily that each "charge" of this expansion correspond to a conserved quantity of the dynamics. The mapping between the boundary conditions and these charges being thus equivalent to the mapping between boundary conditions and constants of the motion. This consideration emphasizes the fact that the deep reason behind the scaling law characterizing the 1d potential free MFG that we consider in this paper is actually their integrable character. A more in depth discussion of this question will appear in a separate publication [15].

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] J.-M. Lasry, P.-L. Lions, C. R. Acad. Sci. Paris, Ser. I 343 (2006) 619.
[2] D. Ullmo, I. Swiecicki, T. Gobron, Phys. Rep. 799 (2019) 1.
[3] A. Lachapelle, J.-M. Lasry, C.-A. Lehalle, P.-L. Lions, Math. Financ. Econ. 10 (2016) 223.
[4] J.-M. Lasry, P.-L. Lions, Jpn. J. Math. 2 (2007) 229.
[5] O. Guéant, J.-M. Lasry, P.-L. Lions, Mean field games and applications, in: Paris-Princeton Lectures on Mathematical Finance 2010, Springer, Heidelberg, 2011, pp. 205–266.
[6] Y. Achdou, F.J. Buera, J.-M. Lasry, P.-L. Lions, B. Moll, Philos. Trans. R. Soc. A 372 (2014) 20130397.
[7] F. Mériauxi, V. Varma, S. Lasaulce, in: 2012 Conference Record of the Forty Sixth Asilomar Conference on Signals, Systems and Computers, ASILOMAR, 2012, pp. 671–675.
[8] A. Kizilkale, R. Malhamé, in: K.G. Vamvoudakis, S. Jagannathan (Eds.), Control of Complex Systems, 2016, pp. 559–584.
[9] D. Bertsekas, Dynamic Programming and Optimal Control, Athena Scientific, Nashua, 2017.
[10] I. Swiecicki, T. Gobron, D. Ullmo, Phys. Rev. Lett. 116 (2016) 128701.
[11] P. Cardaliaguet, J.-M. Lasry, P.-L. Lions, A. Porretta, SIAM J. Control Optim. 51 (2013) 3558, https://doi.org/10.1137/120904184.
[12] F. Dalfovo, S. Giorgini, L.P. Pitaevskii, S. Stringari, Rev. Mod. Phys. 71 (1999) 463.
[13] C.J. Pethick, H. Smith, Bose-Einstein Condensation in Dilute Gases, 2nd ed., Cambridge University Press, 2008.
[14] A. Kamchatnov, Nonlinear Periodic Waves and Their Modulations: An Introductory Course, World Scientific, 2000.
[15] T. Bonnemain, et al., unpublished.