

Second-order PDE's and Deterministic Games

Robert V. Kohn and Sylvia Serfaty*

Abstract. We provide a brief, expository introduction to our recent work on deterministic-game interpretations of second-order parabolic and elliptic PDE.

Mathematics Subject Classification (2000). Primary 49L20, 35J60, 35K55; Secondary 35C99, 49L25, 91A20.

Keywords. deterministic games, motion by curvature, fully nonlinear parabolic equations.

1. Introduction

This note gives a brief, expository introduction to our recent work concerning deterministic game interpretations of some nonlinear second-order PDE's. There are two related but distinct themes:

- *Deterministic control interpretations of geometric evolution laws* [10], and
- *Deterministic control interpretations of fully-nonlinear PDE's* [11].

To capture the main ideas, we shall focus here on simple examples and heuristic arguments. The discussions in [10, 11] are of course quite different – much more general and mathematically rigorous.

2. Deterministic control interpretations of geometric evolution laws

The “level-set method” was introduced in the 1980's as a numerical method for the simulation of geometric evolution laws [12]. Within a few years, it was also recognized as a powerful tool for analyzing the existence and uniqueness of such motions [3, 7].

*The authors gratefully acknowledge support from NSF through grants DMS-0313744 (RVK) and DMS-0239121 (SS), and from the Alfred P. Sloan Foundation through a Sloan Foundation Fellowship (SS).

When the velocity of the moving surface depends only on the normal direction, the level-set description of the motion is a first-order PDE (a Hamilton-Jacobi equation). When the velocity depends on curvature, the level-set description is a second-order parabolic or elliptic PDE. We usually think of first-order and second-order PDE's as being quite different. For example,

- First-order equations have characteristics while second-order parabolic and elliptic equations do not.
- Hamilton-Jacobi equations have well-known links to deterministic control problems (for example, this is the essence of the Hopf-Lax solution formula for $u_t - H(\nabla u) = 0$ when H is convex); for second-order equations the conventional control interpretations are quite different, involving stochastic rather than deterministic control.

The starting point of [10] was the observation that for geometric evolutions, *the first and second-order cases are actually quite similar*. To explain the core idea, let us focus on two key examples:

- (i) *Motion with constant velocity*. Consider the evolution of a region Ω in the plane as its boundary moves inward with constant velocity 1 (Figure 1, left). The evolution is completely characterized by the arrival time

$$u(x) = \text{time that the moving boundary passes through } x.$$

This function solves the stationary Hamilton-Jacobi equation

$$|\nabla u| = 1 \text{ in } \Omega \tag{1}$$

with $u = 0$ at the boundary, and it is characterized by the optimization

$$u(x) = \min_{z \in \partial\Omega} \text{dist}(x, z). \tag{2}$$

- (ii) *Motion by curvature*. Now consider the evolution of a convex region Ω in the plane as its boundary moves with velocity equal to its curvature (Figure 1, right). To track the evolution of the boundary as a parameterized curve we must solve a nonlinear parabolic PDE. But if the region is initially convex then it stays convex, so the evolution is again completely characterized by its arrival time u . A moment's thought reveals that $-\text{div}(\nabla u/|\nabla u|)$ is the curvature of a level set of u , and the velocity of the moving front is $1/|\nabla u|$, so the arrival time of motion by curvature solves

$$|\nabla u| \text{div}(\nabla u/|\nabla u|) + 1 = 0 \text{ in } \Omega \tag{3}$$

with $u = 0$ at the boundary. This PDE is to motion by curvature as the eikonal equation (1) is to motion with constant velocity.




Figure 1. *Left. Motion with constant velocity. Right. Motion by curvature.*

We claim these evolutions are similar in the sense that *motion by curvature also has a deterministic control interpretation*, analogous to (2). It involves a two-person game with players Paul and Carol, and a small parameter ε . Paul is initially at some point $x \in \Omega$; his goal is to exit as soon as possible, and Carol wants to delay his exit as long as possible. The rules are as follows:

- Paul chooses a direction, i.e. a unit vector $|v| = 1$.
- Carol can either accept or reverse Paul's choice, i.e. she chooses $b = \pm 1$.
- Paul then moves distance $\sqrt{2}\varepsilon$ in the possibly-reversed direction, i.e. from x to $x + \sqrt{2}\varepsilon bv$.
- This cycle repeats until Paul reaches $\partial\Omega$.

For example, if Paul is near the top of the rectangle, one might think he should choose v pointing north. But that's a bad idea: if he does so, Carol will reverse him and he'll have to go south (Figure 2, left).




Figure 2. *Left. Paul's quandary – if he tries to go north, Carol will send him south. Middle. Paul can exit from a well-chosen concentric circle in just one step. Right. The construction can be repeated.*

Can Paul exit? Yes indeed. This is easiest to see when $\partial\Omega$ is a circle of radius R . The midpoints of secants of length $2\sqrt{2}\varepsilon$ trace a concentric circle, whose radius is smaller by approximately ε^2/R . Paul can exit in one step if and only if he starts on or outside this concentric circle (Figure 2, middle). This construction can be

repeated of course, producing a sequence of circles from which he can exit in a fixed number of steps (Figure 2, right). Aside from the scale factor of ε^2 they are shrinking with normal velocity $1/R = \text{curvature}$. We have determined Paul's optimal strategy: if $\Omega = B_R(0)$ and his present position is x then his optimal v is perpendicular to x . And we have linked his minimum exit time to motion by curvature.

This calculation is fundamentally local, so it is not really limited to balls. It suggests that Paul's scaled arrival time,

$$u_\varepsilon(x) = \varepsilon^2 \cdot \left\{ \begin{array}{l} \text{minimum number of steps Paul needs to exit starting} \\ \text{from } x, \text{ assuming Carol behaves optimally,} \end{array} \right\} \quad (4)$$

converges as $\varepsilon \rightarrow 0$ to the arrival-time function of motion by curvature. It even provides us with something resembling characteristics for the second-order PDE (3). In fact: Paul's paths are like characteristics, in the sense that the PDE becomes an ODE when restricted to the path (u_ε decreases by exactly ε^2 at each step along Paul's path).

The circle was too easy. How does one analyze more general domains? A key tool is the *dynamic programming principle*:

$$u_\varepsilon(x) = \min_{|v|=1} \max_{b=\pm 1} \left\{ u_\varepsilon(x + \sqrt{2}\varepsilon bv) + \varepsilon^2 \right\}. \quad (5)$$

In words: starting from x , Paul selects the best direction v (taking account that Carol is working against him), recognizing that after taking *this* step he will pursue an optimal path. This principle captures the logic we used in passing from the middle frame of Figure 2 to the right hand frame.

The degenerate-elliptic equation (3) is, in essence, the Hamilton-Jacobi-Bellman equation associated with this dynamic programming principle. To explain why, we use an argument that's familiar from optimal control theory (see e.g. Chapter 10 of [6]). Assume u_ε is smooth enough for Taylor expansion to be valid. Then (5) gives

$$u_\varepsilon(x) \approx \min_{|v|=1} \max_{b=\pm 1} \left\{ u_\varepsilon(x) + \sqrt{2}\varepsilon bv \cdot \nabla u_\varepsilon(x) + \varepsilon^2 \langle D^2 u_\varepsilon(x) v, v \rangle + \varepsilon^2 \right\}, \quad (6)$$

whence

$$0 \approx \min_{|v|=1} \max_{b=\pm 1} \left\{ \sqrt{2}\varepsilon bv \cdot \nabla u_\varepsilon(x) + \varepsilon^2 \langle D^2 u_\varepsilon(x) v, v \rangle + \varepsilon^2 \right\}.$$

Paul should choose v such that $v \cdot \nabla u_\varepsilon(x) = 0$, since otherwise this term will dominate the right hand side and Carol will choose the sign of b to make it positive. In the plane this forces $v = \pm \nabla^\perp u / |\nabla u|$. Either choice is OK: the sign doesn't matter, since the next term is quadratic. We conclude (formally, in the limit $\varepsilon \rightarrow 0$) that

$$\langle D^2 u \frac{\nabla^\perp u}{|\nabla u|}, \frac{\nabla^\perp u}{|\nabla u|} \rangle + 1 = 0.$$

A bit of manipulation reveals that this is the same as (3) in two space dimensions.

To summarize: motion by curvature is similar to motion with constant velocity in the sense that both evolutions can be described by deterministic control problems (the Paul-Carol game versus equation (2)). The PDE that describes the arrival time is, in either case, the associated Hamilton-Jacobi-Bellman equation, derived from the control problem using the principle of dynamic programming. There is, however, a difference: the Paul-Carol game has a small parameter ε , and we only get motion by curvature in the limit $\varepsilon \rightarrow 0$; the optimal control interpretation of the eikonal equation, by contrast, has no small parameter.

This discussion has been formal, and it has focused on just the simplest example. But these ideas can be justified and extended to other geometric motions. In particular:

- The convergence of Paul's scaled arrival time u_ε to the arrival time of motion by curvature can be proved using the framework of "viscosity solutions," following [3, 7]. When u is smooth enough, one can alternatively use a "verification argument;" this gives a stronger result, by estimating the convergence rate. The required smoothness is valid for the arrival time of motion by curvature in the plane; interestingly, however, it fails for the arrival time of motion by mean curvature of a higher-dimensional hypersurface [16].
- The case when Ω is nonconvex is more subtle. Then $\lim_{\varepsilon \rightarrow 0} u_\varepsilon$ is the arrival time of a different motion law, namely the one with normal velocity κ_+ where κ is curvature and $\kappa_+ = \max\{\kappa, 0\}$. The proof depends on a uniqueness result for viscosity solutions, due to Guy Barles and Francesca Da Lio, given in Appendix C of [10].
- These ideas can be extended to higher space dimensions and other geometric evolutions; moreover, the method can be used for parabolic as well as elliptic representations of curvature-driven motion [10]. In addition, a similar approach to some nonlocal geometric evolutions is developed in [9], and a Neumann problem for motion by curvature is addressed in [8].

Our work in this area had important precursors. The Paul-Carol game is essentially a semi-discrete approximation scheme (continuous in space, discrete in time) for motion by curvature. Similar semi-discrete schemes had been considered in the literature on computer vision (e.g. [2, 13, 14]), and in work on numerical schemes for computing viscosity solutions of second-order PDE's [5].

When Paul chooses optimally he becomes indifferent to Carol's choices. One can ask: what happens if Carol just flips a fair coin, but Paul's goal is to arrive *with probability one* in the minimum possible time? A continuous-time version of this problem was studied in [1, 17], as a stochastic-control interpretation of motion by curvature. Paul's optimal choice of direction is same for this stochastic game as in our deterministic setting – roughly speaking, because if he makes a different choice, Carol will take advantage of it with probability 1/2.

By the way, we didn't invent the Paul-Carol game. It was introduced thirty years ago by Joel Spencer, as a heuristic for the study of certain combinatorial problems [18].

3. Deterministic control interpretations of fully nonlinear PDE's

The preceding discussion seems strongly linked to the geometric character of the problem. In particular, Paul's value function u_ε converged to the *level-set description* of a geometric motion. It is natural to whether deterministic game interpretations can also be given for other (non-geometric) second-order PDE's.

The answer is yes! Of course it requires a slightly different perspective. The following deterministic game approach to the 1D linear heat equation was suggested to us by H. Mete Soner. As usual in control theory, we focus on solving a well-posed PDE *backward in time*:

$$v_t + v_{xx} = 0 \text{ for } t < T, \quad \text{with } v = \phi \text{ at } t = T. \quad (7)$$

The associated game has two players; we'll call them Helen and Mark (for a reason to be explained below). There's a marker, that's initially at position $x \in \mathbb{R}$ at time t . At each timestep

- Helen chooses a real number α , then (after hearing Helen's choice) Mark chooses $b = \pm 1$.
- Helen pays penalty $\sqrt{2}\varepsilon\alpha b$.
- The marker moves from x to $x + \sqrt{2}\varepsilon b$ and the clock steps from t to $t + \varepsilon^2$.

The game continues this way until time T . At the final time, Helen collects a bonus $\phi(x(T))$. We referred to Helen's payment of $\sqrt{2}\varepsilon\alpha b$ as a "penalty," but if this number is negative then it actually represents a gain.

We did not yet specify how Helen and Mark make their respective choices. Helen's goal is to maximize her bonus less accumulated penalties. Mark's goal is to give Helen the worst possible result (and Helen knows this). Helen's value function

$$v_\varepsilon(x, t) = \left\{ \begin{array}{l} \text{her optimal final-time bonus minus accumulated penalty,} \\ \text{if the marker starts at position } x \text{ at time } t, \end{array} \right\} \quad (8)$$

is thus determined by the dynamic programming principle

$$v_\varepsilon(x, t) = \max_{\alpha \in \mathbb{R}} \min_{b = \pm 1} \left\{ v_\varepsilon(x + \sqrt{2}\varepsilon b, t + \varepsilon^2) - \sqrt{2}\varepsilon\alpha b \right\} \quad (9)$$

along with the final-time condition $v_\varepsilon(x, T) = \phi(x)$. This leads to the linear heat equation as $\varepsilon \rightarrow 0$ by the same (Taylor-expansion-based) argument used in the last section: proceeding as in (6) and dropping the subscript ε we get

$$v(x, t) \approx \max_{\alpha \in \mathbb{R}} \min_{b = \pm 1} \left\{ v(x, t) + \sqrt{2}\varepsilon b(v_x - \alpha) + \varepsilon^2(v_t + v_{xx}) \right\}. \quad (10)$$

Helen must choose $\alpha = v_x$ to neutralize the term that's linear in ε (otherwise Mark will choose b to make this term negative). This choice of α makes Helen indifferent to Mark's choice of b , so (10) becomes

$$v(x, t) \approx v(x, t) + \varepsilon^2(v_t + v_{xx}).$$

Subtracting $v(x, t)$ from both sides and dividing by ε^2 we conclude that $v_t + v_{xx} = 0$, as desired.

This “game” interpretation of the linear heat equation may seem mysterious, but actually it is rather familiar. In fact it is closely related to the well-known fact that European options can be perfectly hedged in a binomial tree market. In this financial interpretation, with $\varepsilon > 0$,

$$\begin{aligned} x &= \text{the stock price} \\ -\alpha &= \text{the amount of stock in Helen's hedge portfolio} \\ \sum_j \sqrt{2\varepsilon}\alpha_j b_j &= \text{Helen's profit or loss on the hedge portfolio} \\ v_\varepsilon(x, t) &= \text{time-}t \text{ value of the option with payoff } \phi \text{ at time } T, \end{aligned}$$

with the convention that α_j and b_j are Helen's choice of α and Mark's choice of b at time $t_j = t + j\varepsilon^2$. Our players' names come from this interpretation: Helen is the hedger, Mark controls the market. The stock prices are restricted to an (additive) binomial tree, since x increases or decreases by exactly $\sqrt{2\varepsilon}$ at each timestep. The key assertion of perfect hedging is that for a suitable choice of α_j ,

$$v_\varepsilon(x_0, t_0) + \sum_j \sqrt{2\varepsilon}\alpha_j b_j = \phi(x(T)) \quad (11)$$

regardless of how $b_j = \pm 1$ are chosen. Helen is very risk-averse; she always assumes the market (Mark) will move to her detriment. Her optimal α_j are therefore the ones that make (11) true.

The preceding discussion was formal, but (like the arguments in Section 2) it can be fully justified. The rigorous version places a weak upper bound on α (of the form $|\alpha| \leq \varepsilon^{-a}$ where $a > 0$). The main result is that $\lim_{\varepsilon \rightarrow 0} v_\varepsilon(x, t)$ exists and solves the linear heat equation.

Something similar can be done for a large class of fully-nonlinear parabolic and elliptic equations. To explain the main idea, consider a final-value problem of the form

$$v_t + f(Dv, D^2v) = 0 \text{ for } t < T, \quad \text{with } v = \phi \text{ at } t = T \quad (12)$$

on all \mathbb{R}^n . We assume the PDE is (degenerate) parabolic, in the sense that

$$f(p, \Gamma) \leq f(p, \Gamma') \quad \text{if } \Gamma \leq \Gamma' \text{ as symmetric matrices.} \quad (13)$$

The game still has two players (whom we still call Helen and Mark), but the rules are a bit different from before.

1. Helen chooses a vector $p \in \mathbb{R}^n$ and a symmetric $n \times n$ matrix Γ ; then (after hearing Helen's choice) Mark chooses a vector $w \in \mathbb{R}^n$.
2. Helen pays penalty $\varepsilon p \cdot w + \frac{\varepsilon^2}{2} \langle \Gamma w, w \rangle - \varepsilon^2 f(p, \Gamma)$.
3. The marker moves from x to $x + \varepsilon w$ and the clock steps from t to $t + \varepsilon^2$.

The game continues this way until the final time T , when Helen collects a bonus $\phi(x(T))$. Her goal is again to maximize her (worst-case) bonus minus accumulated penalties. Mark does all he can to work against her. Helen's value function $v_\varepsilon(x, t)$ now satisfies the dynamic programming principle

$$v_\varepsilon(x, t) = \max_{p, \Gamma} \min_w \left\{ v_\varepsilon(x + \varepsilon w, t + \varepsilon^2) - \varepsilon p \cdot w - \frac{\varepsilon^2}{2} \langle \Gamma w, w \rangle + \varepsilon^2 f(p, \Gamma) \right\} \quad (14)$$

along with the final-time condition $v_\varepsilon(x, T) = \phi(x)$. To identify (12) as its Hamilton-Jacobi-Bellman equation (in the limit $\varepsilon = 0$) we proceed as usual: using Taylor expansion and dropping the subscript ε , (14) gives

$$v(x, t) \approx \max_{p, \Gamma} \min_w \left\{ v(x, t) + \varepsilon w \cdot (\nabla v - p) + \varepsilon^2 \left(\frac{1}{2} \langle (D^2 v - \Gamma) w, w \rangle + f(p, \Gamma) + v_t \right) \right\}. \quad (15)$$

Helen must choose $p = \nabla v$ to neutralize the term that's linear in ε (otherwise Mark will choose w to make this term dominant, working against her). She also needs $\Gamma \leq D^2 v$ (otherwise Mark can drive $\langle (D^2 v - \Gamma) w, w \rangle$ to $-\infty$ by a suitable choice of w). For such p and Γ , the right hand side of (15) reduces to

$$\max_{\Gamma \leq D^2 v} \left\{ v(x, t) + \varepsilon^2 (f(Dv, \Gamma) + v_t) \right\}$$

The optimal Γ is $D^2 v$, as a consequence of the parabolicity condition (13), leading as asserted to the formal HJB equation $v_t + f(Dv, D^2 v) = 0$.

For the linear heat equation (7) Helen had only to choose $\alpha \in \mathbb{R}$. For the fully-nonlinear equation (12) she had to choose *both* a vector p *and* a matrix Γ . Reviewing the arguments, we see why. When the equation is nonlinear, we need separate proxies for Dv and $D^2 v$. The vector p is a proxy for the former, while the matrix Γ is a proxy for the latter.

The calculations presented here are of course purely formal. The solution of a fully-nonlinear PDE like (12) need not be smooth, nor even C^1 . The rigorous analysis uses viscosity-solution methods, showing that the "semi-relaxed limits"

$$\begin{aligned} \bar{v}(x, t) &= \limsup_{y \rightarrow x, s \rightarrow t, \varepsilon \rightarrow 0} v_\varepsilon(y, s) \\ \underline{v}(x, t) &= \liminf_{y \rightarrow x, s \rightarrow t, \varepsilon \rightarrow 0} v_\varepsilon(y, s) \end{aligned}$$

are respectively a subsolution and a supersolution of the PDE. If f is such that the PDE has a comparison principle, then it follows that $\bar{v} = \underline{v}$.

Our two-person game for the linear heat equation was related to hedging in a binomial market. It can thus be viewed as a discrete-time, deterministic version of the Black-Scholes approach to option pricing. Our game for the fully-nonlinear parabolic equation (12) is, in a similar sense, a discrete-time, deterministic version of the stochastic representation formula developed in [4].

4. Discussion

The following table encapsulates some well-known connections between PDE's and applications. Hamilton-Jacobi equations frequently come from optimal control

problems. The linear heat equation is steepest descent for the Dirichlet integral, and motion by curvature is steepest descent for perimeter. The linear heat equation is also linked to Brownian motion, and the value function of a stochastic control problem solves a nonlinear parabolic PDE.

	deterministic control	stochastic control	steepest descent
1st order nonlinear	Hamilton-Jacobi		
2nd order linear	***	random walk	$\int \nabla u ^2$
2nd order nonlinear	***	controlled diffusion	perimeter

Our contribution has been to add two additional connections, corresponding to the asterisks in the table:

- (1) We have shown that motion by curvature has a deterministic control interpretation; indeed, its level-set representation is roughly speaking the Hamilton-Jacobi-Bellman equation of a two-person game. Our discussion focused for simplicity mainly on the motion of convex curves in the plane, but the viewpoint is much more general.
- (2) We have shown that many nonlinear PDE's have deterministic control interpretations. The main requirement is that the PDE have a comparison principle (and therefore a unique viscosity solution). When restricted to the linear heat equation, our interpretation is closely connected to the pricing and hedging of options in a binomial tree market.

These connections are, we think, of intrinsic interest. Perhaps they may also have practical value. We close with two questions about possible directions for further work.

- *Can our deterministic control interpretations be used to prove new results about PDE?* Here the games in Section 2 seem more promising. In fact, our paper [10] includes a modest application of this type: a “waiting-time” result for motion with velocity κ_+ (Theorem 7). The games in Section 3 seem less promising, because they are virtually equation-independent. Of course, if the goal is to derive new PDE results, there is no reason to restrict attention to *deterministic* games. The recent paper [15] provides a fine example of how an equation-dependent (but stochastic) control interpretation can be used to derive new results about a nonlinear PDE (namely the “infinity-Laplacian”).

- *Can our deterministic control interpretation be the basis of a numerical solution scheme?* As noted in Section 2, our interpretation of motion by curvature is closely connected to the numerical solution schemes for curvature-driven motion studied in [2, 5, 13, 14]. Concerning Section 3: the dynamic programming principle (14) amounts to a semidiscrete time-stepping scheme for (12). When the solution is smooth it amounts to explicit Euler, since the optimal p and Γ are Dv and D^2v respectively. So (14) is a version of explicit Euler that works even if the solution is not C^1 . Can this time-stepping scheme be approximated numerically in a spatially discrete setting?

Acknowledgement. Much of Section 2 appeared previously in an article “Parabolic PDE’s and Deterministic Games,” by Robert V. Kohn, *SIAM News* **40**, no. 8 (Oct. 2007).

References

- [1] Buckdahn, R., Caraliaguet, P., and Quincampoix, M., A representation formula for the mean curvature motion, *SIAM J. Math. Anal.* **33** (2001) 827–846
- [2] Catté, F., Dibos, F., and Koepfler, G., A morphological scheme for mean curvature motion and applications to anisotropic diffusion and motion of level sets, *SIAM J. Numer. Anal.* **32** (1995) 1895–1909
- [3] Chen, Y.G., Giga, Y., and Goto, S., Uniqueness and existence of viscosity solutions of generalized mean curvature flow equations, *J. Diff. Geom.* **33** (1991) 749–786
- [4] Cheridito, P., Soner, H.M., Touzi, N., and Victoir, N., Second-order backward stochastic differential equations and fully nonlinear parabolic PDE’s, *Comm. Pure Appl. Math.* **60** (2006) 1081–1110
- [5] Crandall, M. and Lions, P.L., Convergent difference schemes for nonlinear parabolic equations and mean curvature motion, *Numer. Math.* **75** (1996) 17–41
- [6] Evans, L.C., *Partial Differential Equations*, American Mathematical Society (1997)
- [7] Evans, L.C. and Spruck, J., Motion of level sets by mean curvature I, *J. Diff. Geom.* **33** (1991) 635–681
- [8] Giga, Y. and Liu, Q., in preparation.
- [9] Imbert, C. and Serfaty, S., in preparation.
- [10] Kohn, R.V. and Serfaty, S. (with an appendix by G. Barles and F. Da Lio), A deterministic-control-based approach to motion by curvature, *Comm. Pure Appl. Math.* **59** (2006) 344–407
- [11] Kohn, R.V. and Serfaty, S. in preparation.
- [12] Osher, S. and Sethian, J.A., Fronts propagating with curvature-dependent speed – Algorithms based on Hamilton-Jacobi formulations, *J. Comp. Phys.* **79** (1988) 12–49
- [13] Guichard, F., Thèse, Université de Paris – Dauphine, Paris (1994)
- [14] Pasquignon, D., Approximation of viscosity solutions by morphological filters, *ESAIM: Control Optim. and Calc. of Var.* **4** (1999) 335–359

- [15] Peres, Y., Schramm, O., Sheffield, S., and Wilson, D., Tug of war and the infinity Laplacian, *J. Amer. Math. Soc.*, in press.
- [16] Sesum, N., Rate of convergence of the mean curvature flow, *Comm. Pure Appl. Math.* **61** (2007) 464–485
- [17] Soner, H.M. and Touzi, N., A stochastic representation for mean curvature type geometric flows, *Ann. Prob.* **31** (2003) 1145–1165
- [18] Spencer, J., Balancing games, *J. Combinatorial Th. Ser B* **23** (1977) 68–74

Robert V. Kohn, Courant Institute of Mathematical Sciences, 251 Mercer Street, New York, NY 10012

E-mail: kohn@cims.nyu.edu

Sylvia Serfaty, Courant Institute of Mathematical Sciences, 251 Mercer Street, New York, NY 10012

E-mail: serfaty@cims.nyu.edu