Model reduction for agent-based social simulation: Coarse-graining a civil violence model

Yu Zou,[1] Vladimir A. Fonoberov,[2] Maria Fonoberova,[2] Igor Mezic,[2,3] and Ioannis G. Kevrekidis[1,*]

[1]*Department of Chemical and Biological Engineering and Program in Applied and Computational Mathematics, Princeton University, Princeton, New Jersey 08544, USA*

[2]*Aimdyn, Inc., 1919 State Street, Suite 207, Santa Barbara, California 93101, USA*

[3]*Department of Mechanical Engineering, University of California, Santa Barbara, Santa Barbara, California 93106, USA*

(Received 23 April 2011; published 8 June 2012)

Agent-based modeling (ABM) constitutes a powerful computational tool for the exploration of phenomena involving emergent dynamic behavior in the social sciences. This paper demonstrates a computer-assisted approach that bridges the significant gap between the single-agent microscopic level and the macroscopic (coarse-grained population) level, where fundamental questions must be rationally answered and policies guiding the emergent dynamics devised. Our approach will be illustrated through an agent-based model of civil violence. This spatiotemporally varying ABM incorporates interactions between a heterogeneous population of citizens [active (insurgent), inactive, or jailed] and a population of police officers. Detailed simulations exhibit an equilibrium punctuated by periods of social upheavals. We show how to effectively reduce the agent-based dynamics to a stochastic model with only two coarse-grained degrees of freedom: the number of jailed citizens and the number of active ones. The coarse-grained model captures the ABM dynamics while drastically reducing the computation time (by a factor of approximately 20).

DOI: 10.1103/PhysRevE.85.066106 PACS number(s): 89.65.−s, 89.75.−k, 05.10.Gg

I. INTRODUCTION

Agent-based modeling (ABM) has become a powerful tool for studying the behavior of complex systems in recent years in part through advances in modern computing technology. In this modeling approach the system studied consists of individual agents whose actions and interactions are computed via stipulated behavioral rules. The ABM provides a convenient framework in which different interactions between individuals in possibly large, heterogeneous populations can be computationally implemented and the dynamics and phenomena resulting from these interaction rules can be observed and hopefully rationalized. The use of ABM has a long lineage dating back to von Neumann's self-reproducing automata [1]. Pioneered by Schelling through his seminal works [2–4], the ABM has found increasing application in studying a wide variety of topics in social science (see, e.g., Refs. [5–8]).

Agent-based modeling has, however, an inherent weakness in that the computational time for large complex systems, which typically consist of thousands or millions of agents, can be prohibitively long. This drawback makes the comprehensive analysis of certain ABM results extremely time consuming and, as a result, may unsatisfactorily delay realistic decision and policy making associated with important societal, economical, or military events. Improving computational efficiency for ABM therefore becomes a necessity. This naturally links with mathematical approaches to coarse graining and model reduction, techniques that attempt to single out crucial macroscopic factors of complex systems and construct accurate but reduced alternatives to the original full-system dynamics. In this manner the extraction of useful model predictions can be significantly accelerated.

One approach to achieving an accurate and efficient model reduction that is particularly well suited for ABM is the recently developed equation-free (EF) framework [9–11]. This framework circumvents the explicit derivation of macroscopic, system-level equations for coarse-grained observables (statistics and features) of the ABM; it is applicable when we believe that such coarse-grained equations in principle exist, yet they are not available in closed form. Our purpose in this paper is to show how to implement the EF method to ABM, obtaining reduced models and significantly accelerating the computational extraction of system-level information from the ABM.

The concrete topic we choose to illustrate this application is an agent-based civil violence model originally proposed by Epstein [12]. In that paper insurgency was characterized by repeated outbursts of active citizen populations, arising in a simulation in the form of a punctuated equilibrium, a phenomenon widely observed across disciplines. From the perspective of dynamics, it would be interesting to model this punctuated equilibrium behavior via a few representative coarse variables. Furthermore, if such a reduced model can be obtained, the simulation time for the original civil violence model may also decrease significantly. Since the dynamics of any system in punctuated equilibrium is stochastic, we attempt to build the reduced effective model in the form of a system of stochastic differential equations (SDEs) in which one or more terms are a stochastic process. The resulting solution is then in itself a stochastic process [13]. A key difficulty in constructing the reduced model as a multidimensional (here two-dimensional) SDE lies in accurately obtaining the drift and diffusion coefficients in the SDE as functions of the coarse variables. As we will discuss in more detail below, we will use two characteristic aggregate features of the original model—the numbers of jailed and active citizens—as our two coarse variables. We will also design a computational procedure, i.e., the so-called lifting step in the EF framework, that generates detailed agent states consistent with prescribed values of the coarse variables, thus enabling an estimation of the drift and diffusion coefficients. The particular lifting

*yannis@princeton.edu

 ©2012 American Physical Society

procedure we will use is inspired from observations of the dynamics of the detailed ABM.

The paper is organized as follows. Section II briefly describes the original civil violence model derived from Epstein [12] and showcases its characteristic dynamical behavior. Section III provides our observations on relationships between our two chosen coarse variables and the remaining system variables. We then design the lifting procedure and show how to estimate the drift and diffusion coefficients of the reduced model through short bursts of appropriately designed ABM simulation. Significant statistics of the temporal behavior obtained through both the original and the reduced model will be compared as well. In Sec. IV the Fokker-Planck equation and the backward Kolmogorov equation, based on our two-dimensional reduced effective SDE, will be solved to obtain, respectively, the long-term stationary joint probability density function (PDF) of the two coarse variables and the mean exit time for the system (described by the two coarse variables) to escape a given domain (leading to insurgency). This mean exit time can be used to approximate the average time between social upheavals. A summary will be presented in Sec. V.

II. CIVIL VIOLENCE ABM AND SIMULATION

In Epstein's original model [12] the dynamics of social violence is simulated via interactions among citizens and police officers that are placed on a square lattice grid. Each position in the lattice grid can be occupied by one agent (citizen or police officer) only. A citizen is defined as an agent with two inherent characteristics: hardship H and risk aversion R. The values of these characteristics are sampled from the uniform distribution $U(0,1)$ and fixed for each agent during the entire simulation. Quantifying the perceived legitimacy of the regime by a quantity L, a citizen will decide whether or not to rebel based on the following rule: If $H(1-L) - RP > T$, then the citizen becomes active; otherwise, the citizen remains inactive (quiet). In the above rule T is a (fixed) threshold value and P is a (spatiotemporally varying) arrest probability that depends on the current state of the agent neighborhood.

Active citizens may be arrested, and jailed, by the police officers in their spatial vicinity. The arrest probability is modeled in Ref. [12] as

$$P = 1 - \exp[-k(C/A)_v], \tag{1}$$

where C and A are, respectively, the numbers of police officers and active citizens in a circular neighborhood with radius v (called vision v) and k is a constant. In this paper a different model for the arrest probability, proposed in Ref. [14], is used:

$$P = 1 - \exp[-k(C/A)_v] \sum_{i=0}^{15} \frac{[k(C/A)_v]^i}{i!}. \tag{2}$$

The sigmoidal function given by Eq. (2) is capable of giving rise to insurgency outbursts (see Ref. [14] for the justification of this modification to the Epstein model and for references documenting the occurrence of the punctuated equilibrium phenomenon in this social context). The shapes of the arrest probability from the above two models are compared in Fig. 1.

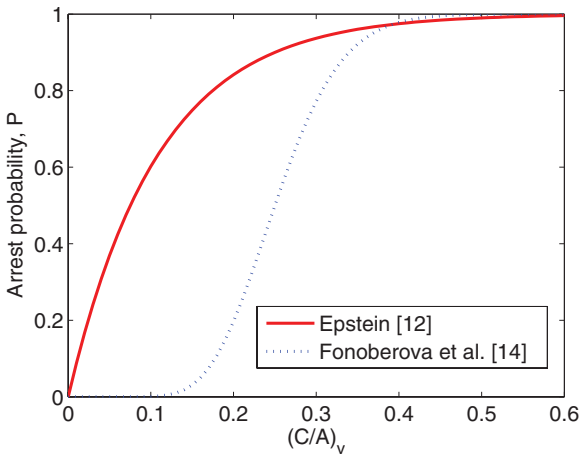

FIG. 1. (Color online) Arrest probability from Eqs. (1) and (2) (see the text).

Citizens and police officers constantly perform random walks on the lattice by moving to a randomly selected position in their size-1 Moore neighborhood if that position is not occupied. Otherwise, they remain in their original position. Police officers also have the ability to arrest the nearest active citizen in their vision v', jump to that citizen's position, and send the citizen to jail. The jail term of an arrested citizen is variable; at every arrest it is initialized with a random integer sampled over the uniform distribution $U(0, J_{\max})$. As citizens complete their full term in the jail, they are released back to the lattice and select random, unoccupied locations to reside.

Within each day (one step of the evolution algorithm is taken to be one day), a citizen is allowed to move only once, yet the police officers are more mobile and can move and arrest M (>1) times (this is another modification of the Epstein model proposed in Ref. [14]). Table I lists values of the parameters used to simulate the model. Note that the parameter k is found from the condition that $P = 0.5$ when $M = A/C$. The boundary conditions in space are set to be periodic, which means that an agent, after exiting the domain at one of the four side boundaries, reenters the lattice from the opposite side.

A typical simulation procedure within a one-day step is described by the flow chart in Fig. 2. Using the model and the procedure described above, we perform multiagent simulations and extract time histories of the numbers of jailed

TABLE I. Parameters used by the simulation (for notation see the text).

Parameter	Value
lattice size	100×100
citizen density	0.7
police officer density	0.01
T	0.1
L	0.8
$J_{\max}$	120
k	62.6716
v	14
v'	14
M	4

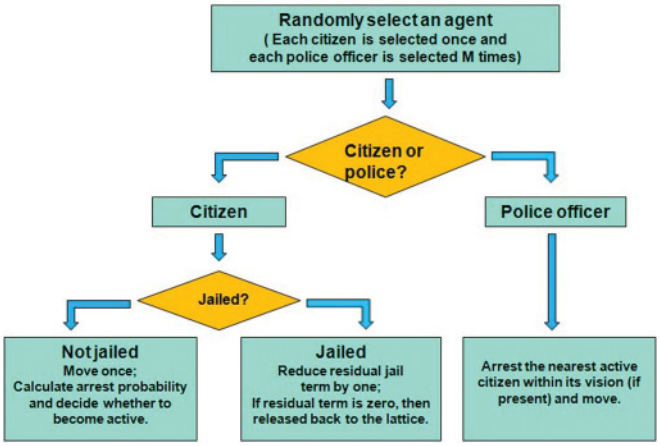

FIG. 2. (Color online) Flow chart for the model simulation procedure within a one-day period.

and active citizens, as plotted in Fig. 3(a). The histories show remarkable emergent dynamics exhibiting the so-called punctuated equilibrium phenomenon [15,16]: Long periods of quiescent, stable stasis are punctuated by almost instantaneous outbursts of insurgency. Note, however, that this phenomenon does not occur in the original Epstein model [12] with the arrest probability P defined as in Eq. (1) as discussed in Ref. [14]. The observations are also plotted in a two-dimensional phase plane projection [Fig. 3(b)].

III. COARSE GRAINING: SDE REDUCED MODEL APPROXIMATION

The punctuated equilibrium shown in Fig. 3 is a feature of the collective behavior of the citizen-police system whose occurrence results from detailed interactions among all agents. Table II lists inherent (fixed during a simulation) properties as well as (time-varying) state variables for each type of agent. The question arises naturally as to whether it is possible to obtain a closed set of dynamic equations for the evolution of only a few coarse variables. More specifically we ask if such a set of dynamic equations can be written in terms of only two such variables: the numbers of the jailed and the active citizens. In other words, if the information we have is the current value of only these two coarse variables, we query if it is possible to obtain a model that allows us to predict the value of these two variables in the future. To date, no explicit set of equations, deterministic [ordinary differential equations (ODEs)] or stochastic [stochastic differential equations (SDEs)], has been obtained analytically that accurately describes the temporal evolution in terms of only these two quantities.

If such a set of equations can in principle be derived, it makes sense to expect that this reduced closed model will be stochastic in the form of a SDE

$$\begin{aligned} dX_1 &= \mu_1 dt + \sigma_1 dW_1, \\ dX_2 &= \mu_2 dt + \sigma_{21} dW_1 + \sigma_{22} dW_2, \end{aligned} \tag{3}$$

where X_1 and X_2 are, respectively, the numbers of jailed and active citizens and W_1 and W_2 are independent Wiener processes. Then the drift and diffusion coefficients $\mu_1,\mu_2,\sigma_1,\sigma_{21}$, and σ_{22} will, in general, be functions of X_1 and X_2 only. To evaluate these functions for different values of X_1 and X_2, the following equation can be used (see, e.g., Ref. [17]):

$$\begin{aligned} \mu_i(x_1,x_2) &= \lim_{\Delta t\to 0} \frac{\langle X_i(t+\Delta t) - X_i(t)\rangle}{\Delta t}\bigg|_{X_1(t)=x_1, X_2(t)=x_2}, \\ D_{ij}(x_1,x_2) &= \lim_{\Delta t\to 0} \frac{\langle X_i(t+\Delta t) - X_i(t) - \mu_i \Delta t\rangle\langle X_j(t+\Delta t) - X_j(t) - \mu_j \Delta t\rangle}{\Delta t}\bigg|_{X_1(t)=x_1, X_2(t)=x_2}. \end{aligned} \tag{4}$$

The operator $\langle\cdot\rangle$ represents ensemble average and D_{ij} are the diffusion coefficients of the Fokker-Planck equation corresponding to the SDE (3). In this paper the values of μ_i and D_{ij} are estimated using

$$\begin{aligned} \mu_i(x_1,x_2) &\approx \frac{\langle X_i(t+\Delta t) - X_i(t)\rangle}{\Delta t}\bigg|_{X_1(t)=x_1, X_2(t)=x_2}, \\ D_{ij}(x_1,x_2) &\approx \frac{\langle X_i(t+\Delta t) - X_i(t) - \mu_i \Delta t\rangle\langle X_j(t+\Delta t) - X_j(t) - \mu_j \Delta t\rangle}{\Delta t}\bigg|_{X_1(t)=x_1, X_2(t)=x_2} \end{aligned} \tag{5}$$

by setting Δt to a single day ($\Delta t = 1$): The diffusion coefficients in Eq. (3) are then calculated via a Cholesky decomposition [17,18]

$$\sigma_1 = \sqrt{D_{11}}, \quad \sigma_{21} = \frac{D_{12}}{\sigma_1}, \quad \sigma_{22} = \sqrt{D_{22} - \sigma_{21}^2}. \tag{6}$$

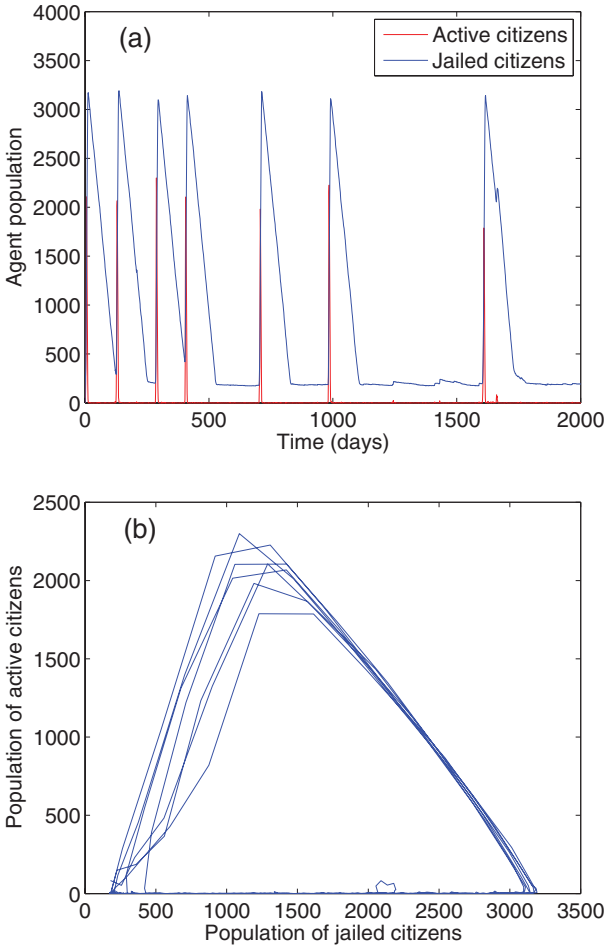

FIG. 3. (Color) (a) Time history of the numbers of jailed and active citizens extracted from the full agent-based simulation and (b) their two-dimensional phase plane projection.

In Eq. (5) we need to compute the ensemble values of $X_i(t+\Delta t) - X_i(t)$, with $i = 1,2$, for (in principle) any possible pair $[X_1(t), X_2(t)] = (x_1, x_2)$. This can be achieved via a properly designed coarse time stepper for X_i, with $i = 1,2$. Starting with a pair of values for $X_1(t)$ and $X_2(t)$, we need to be able to generate consistent detailed initial conditions for the entire population of all agents (this is called the lifting procedure in the equation-free framework). The entire system then evolves over a short time interval Δt via the agent-based simulation and the values of the coarse variables (the numbers of jailed and active citizens) $X_i(t+\Delta t)$ are observed from the detailed simulation; this is the so-called restriction step. Clearly, the restriction step is easy: At the end of the full ABM simulation we simply count the numbers of jailed and active citizens; the key difficulty lies in finding an appropriate lifting procedure. In what follows a detailed approach to implementing a reasonable lifting is presented and discussed.

A. From macro to micro: Discussion of lifting considerations

Suppose that a simulation is suddenly interrupted and from it we are given only the values (at that time) of the two coarse variables: the numbers of jailed X_1 and active citizens X_2 (and we also know the constant during the entire simulation, i.e., the total number of citizens N_{citizen}). A very interesting and important question then becomes how one can construct appropriate realizations of the entire detailed ABM state (lattices of active as well as inactive citizens and police officers, jailed citizens, and their jail term distributions) that are consistent with the two coarse values. Appropriate here means that whether we continue the original, interrupted simulation or restart it from these artificially constructed (lifted) states, the ABM dynamics will be effectively the same (in probability). It might be tempting, just knowing X_1 and X_2, to initialize all Table II features for all agents in a random and uncorrelated way; yet we find that in order to reproduce the true system dynamics certain correlations that have developed during the simulation must somehow be captured and retained when consistently initializing the detailed ABM states.

It is worth pausing for a moment to consider how much information must be reinjected, in this step, in the fine scale model: Based on only two scalar numbers we must assign hundreds and even thousands of agents (citizens and policemen) with appropriate intrinsic properties on the lattice as well as citizens in jail with appropriate jail terms. In certain cases, if there is a large separation of time scales in the problem, any errors one makes in these assignments are quickly forgotten after a few simulation steps. In our problem this is not the case: In addition to the values of the coarse variables we will incorporate certain qualitative observations we have made, based on extensive ABM simulations, in our construction of an acceptable lifting operator.

TABLE II. Inherent properties and state variables for the different types of agents (see the text). A check mark indicates that this type of agent has a given property; the dash indicates that it does not.

Agent type	H	R	Position	Active status	Residual jail term
active citizens	√	√	√	active	—
inactive citizens	√	√	√	inactive	—
jailed citizens	√	√	—	inactive	√
police officers	—	—	√	—	—

(a)

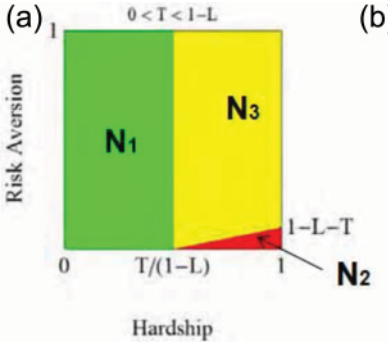

(b)

Citizen type \ Region	Region 1 green	Region 2 red	Region 3 yellow
Active citizen	N/A	A_2	X_2-A_2
Inactive citizen	N_1	N/A	$N_{citizen}$-X_1-X_2-N_1
Jailed citizen	N/A	N_2-A_2	X_1-(N_2-A_2)

FIG. 4. (Color online) (a) Partition of the H-R domain into three regions and (b) allocation of H,R pairs in each region to each citizen type (active, inactive, and jailed).

We start by observing that the two-dimensional parameter space of all citizens' H and R values (which, as intrinsic properties of the agents, are prescribed at the beginning of the simulation and remain constant through it) can be divided into three parts [Fig. 4(a)].

Region 1: $H \leqslant T/(1-L)$. Citizens in this region are always inactive.

Region 2: $H > T/(1-L)$ and $H(1-L) - R > T$. Citizens in this region are always active or jailed.

Region 3: $H > T/(1-L)$ and $H(1-L) - R \leqslant T$. Citizens in this region may be active, inactive, or jailed. Whether they are active or inactive depends on the value of the arrest probability P, which varies in space and time.

With the denotation of the number of citizens in region 1 of the H-R parameter plane by N_i (these are fixed at the beginning of the simulation), the allocation of subpopulation types (active, inactive, or jailed) in each of the three regions is shown in the chart of Fig. 4(b). There is only one unknown quantity in these allocations: the value of A_2, which is constrained by $A_2 < X_2$ and $N_2 - A_2 < X_1$. At the beginning of every computational experiment, given the current values of the coarse variables X_1 and X_2, we first select the value of A_2 at random. Given the prescriptions in Fig. 4(b), we calculate the subpopulation count in each region. The available H-R pairs are then randomly assigned to citizens based on the allocations indicated in the chart. For example, all N_1 pairs in region 1 are allocated to inactive citizens; N_2-A_2 pairs are randomly selected from region 2 and allocated to jailed citizens; X_2-A_2 pairs are randomly selected from region 3 and allocated to active citizens, and so on. Note that, because citizens in region 2 are always active or jailed, we have $X_1 + X_2 \geqslant N_2$. In addition, since all citizens in region 1 never turn active or go to jail, we have $X_1 + X_2 \leqslant N_{\text{citizen}} - N_1 = N_2 + N_3$. Also accounting for the conditions $X_1 \geqslant 0$ and $X_2 \geqslant 0$, the trajectories in the X_1-X_2 phase plane are always confined in the simulation domain as shown in Fig. 5.

The police officer positions change drastically during an outburst of rebellious activity; during the quiescent period, however, when the population of active citizens is very close to zero, the police officer positions show more (if not entirely) homogeneous patterns. This can be seen in Figs. 6(b)–6(k), where we divide the simulation lattice into a 10×10 grid and plot populations of police officers in each grid cell.

The spatial heterogeneity of police officer position patterns is evaluated via the following measure:

$$\sigma = \frac{\sum_{i,j=1}^{10} (C_{(i,j)} - C_{\text{hom},(i,j)})^2}{100}, \tag{7}$$

where $C_{(i,j)}$ is the number of police officers in cell (i,j) and $C_{\text{hom},(i,j)}$ is the number of police officers in cell (i,j) if the police officers were homogeneously distributed. Clearly, since there are 100 total 10×10 coarse cells in our 100×100 lattice and a total of 100 police officers, $C_{\text{hom},(i,j)} = 1$ for each of these coarse cells (i,j), with $i,j = 1,\ldots,10$. The time history of the heterogeneity measure σ is shown in Fig. 6(l).

The reason for the above temporal variation of the police officer position patterns is that police officers move rapidly to arrest active citizens in an outburst, but after the outburst is suppressed, there are almost no citizens to arrest, so the police officers move only in their size-1 Moore neighborhood. Their position patterns therefore become nearly stationary between two successive outbursts. As a culmination of the quick outbursts of rebellions, these position patterns of police officers appear to be able to effectively suppress almost all possible subsequent insurgencies; we observe that they maximize arrest

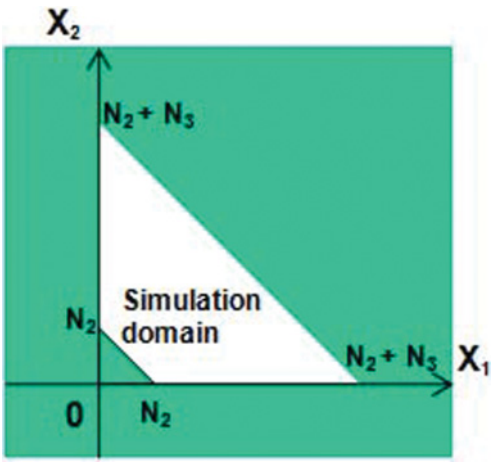

FIG. 5. (Color online) Simulation domain in the X_1-X_2 phase plane that contains all trajectories. The central white region stands for the simulation domain. The surrounding green region is not accessible by X_1 and X_2 pairs.

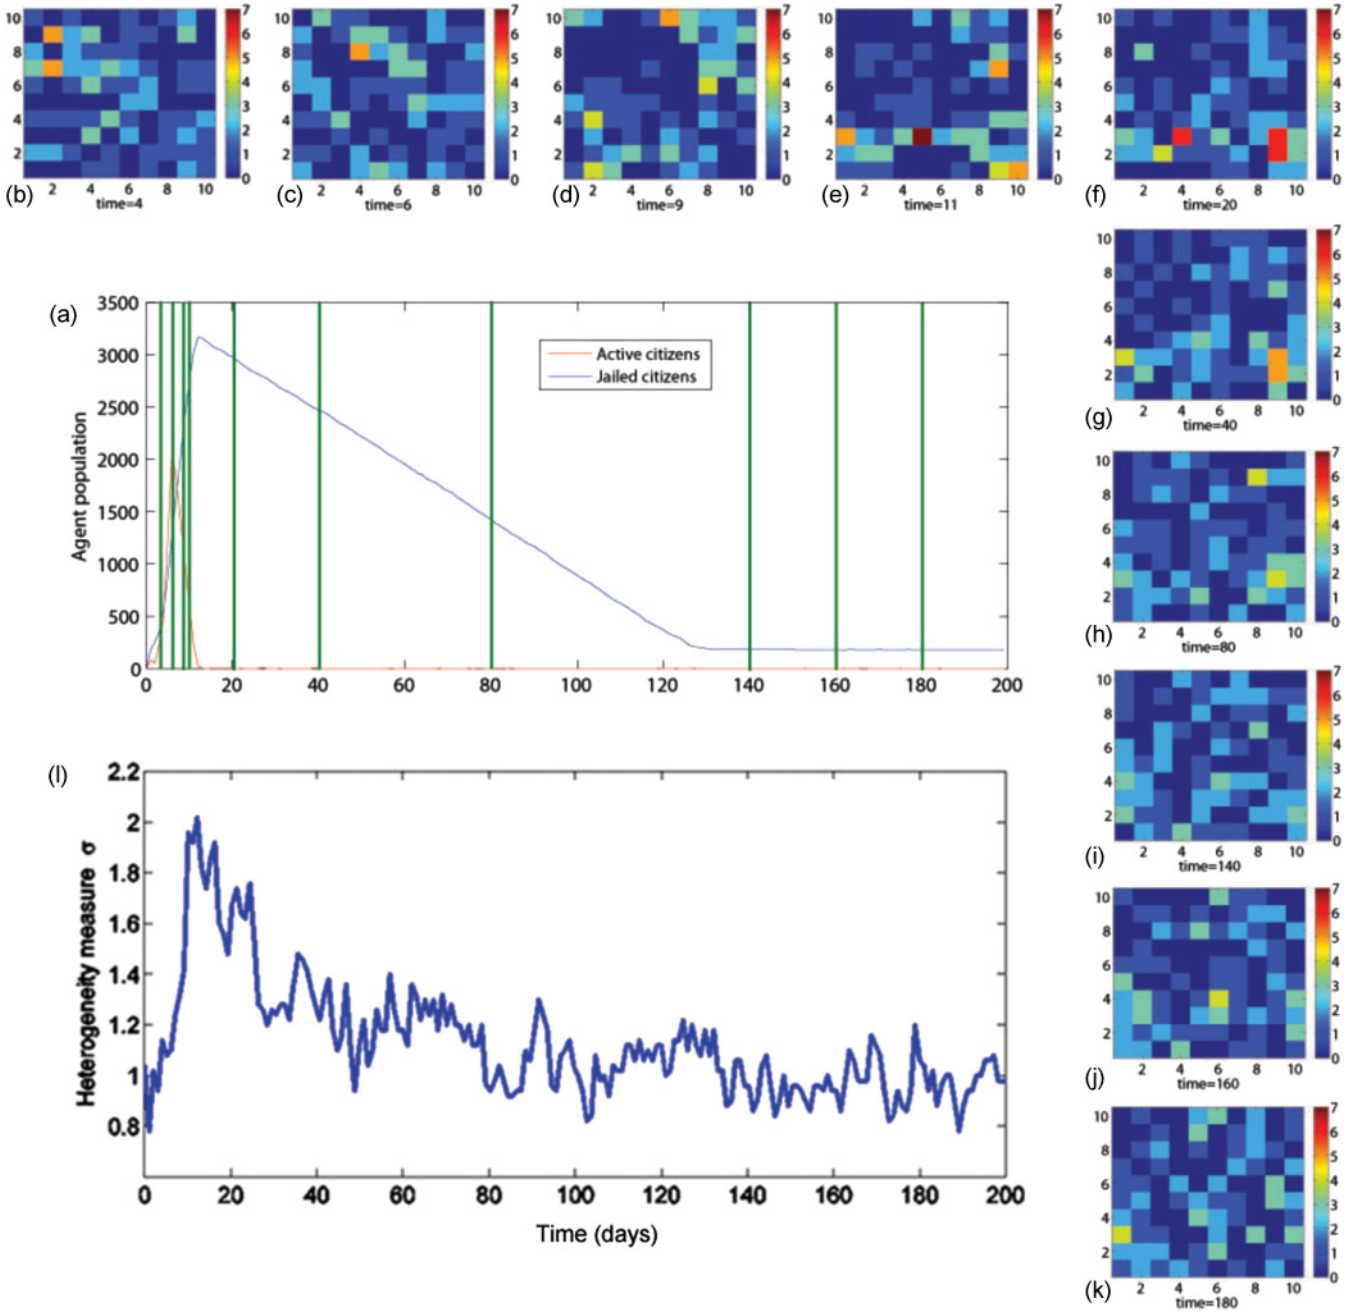

FIG. 6. (Color) Snapshots and heterogeneity measures of police officer position patterns between time 0 and 200 in a single realization of simulation. (a) Time history for populations of jailed and active citizens. Green vertical bars stand for the time steps at which the police officer position patterns shown have been obtained. (b)–(k) Snapshots of police officer position patterns in a 10×10 grid of uniform cells. (l) Heterogeneity measure σ of the police officer position patterns. Drastic police officer movement can be seen between times 4 and 20. More homogeneous and stationary position patterns subsequently arise between times 20 and 200.

probabilities of potentially active citizens in region 3, leaving only a small chance for an outburst to occur again. This is also the cause for a long stable stasis between two adjacent rebellions. To consistently generate police officer positions, one possible way is to run the original simulation for some time and record one snapshot of police officer positions within the quiescent period where the number of active citizens X_2 is low, say, 5–10. This set of positions is then used to place police officers for (X_1, X_2) pairs with X_2 in a representative range, say $[0,8]$. Police officer positions are generated randomly for X_2 values out of this range.

By monitoring several snapshots of the citizen positions during an outburst, we observe that active citizens are almost invariably located fairly homogeneously within a circle whose radius r varies with time (see Fig. 7). We also observe that the density of active citizens within this circle is (to a very good approximation) equal to the density of potentially active citizens [that is, citizens with internal properties H and R in

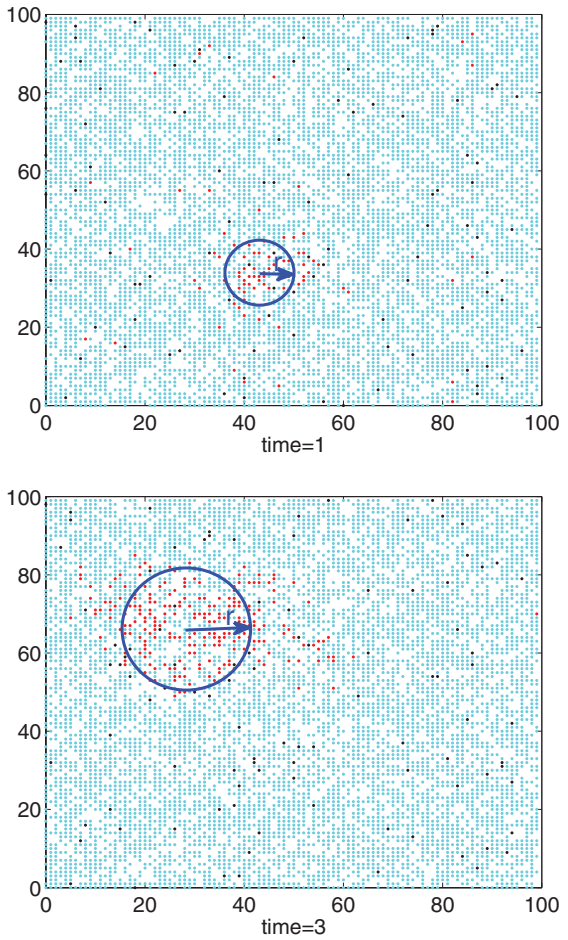

FIG. 7. (Color) Two snapshots of active citizen position patterns at times 1 and 3 in the simulation of Fig. 6. Red, cyan, and black dots represent active citizens, inactive citizens, and police officers, respectively. The active citizen positions can be approximately enclosed by a circle with the radius r calculated from Eq. (8).

region 2 and 3 of Fig. 4(a)] over the entire lattice. Therefore, the value of $r(t)$ can be estimated by

$$r(t) \approx \sqrt{\frac{X_2(t)}{\pi(N_2+N_3)/L^2}}. \tag{8}$$

When initializing citizen positions, given $X_2(t)$, we therefore place active citizens uniformly within a circle with the appropriate radius r. Since the boundary conditions are periodic in space, the center of the circle can be set (without loss of generality) at the center of the lattice. We then place all other inactive and nonjailed citizens, with properties H and R in regions 2 and 3, randomly and uniformly outside this circle. Finally all citizens with properties H and R in region 1 are randomly and uniformly placed on the square lattice.

We also observe that the residual jail terms for jailed citizens show interesting patterns. If the distribution of residual jail terms is plotted separately for jailed citizens with properties in region 2 and those with properties in region 3, we observe that the distribution for region 2 jailed citizens is almost always uniform over the domain $[0, J_{\max}]$ [Figs. 8(a) and 8(b)]. The distribution for region 3 jailed citizens, in contrast, is nearly uniform, with a density of about 25 for the parameter set (Table I) used in this paper [Figs. 8(c) and 8(d)]. The upper limit of this latter distribution can thus be estimated from the population of region 3 jailed citizens, i.e., $X_1 - (N_2 - A_2)$ in Fig. 4(b).

B. Simulating the reduced model

Based on the above lifting procedure, we can construct initial conditions for the entire system with the prescribed distribution of inherent properties H and R and with the state variables (positions, active status, and residual jail terms) of all agents consistently initialized based on the numbers of jailed and active citizens only. We can then proceed to estimate the drift and diffusion coefficients in Eq. (3) via brief computational bursts of ensemble realizations of the original ABM and the formulas in Eqs. (4) and (6). The ensemble size for performing the ensemble average in Eq. (4) is set to 200. We estimate these coefficients as functions of X_1 and X_2 on a two-dimensional grid with a spacing of 1 within the domain $[0,300] \times [0,20]$ and a spacing of 10 otherwise. Computations were not performed in the region $X_1 + X_2 < N_2$ since X_1 and X_2 cannot actively enter that region. The coefficients for X_1 and X_2 values not on the grid are interpolated from the closest three grid values through a linear combination of linear (finite-element basis) shape functions.

We have now obtained our effective reduced SDE (3); to simulate it we also need boundary conditions in X_1-X_2 space, which are set as follows:

$$\begin{aligned} &\text{if } X_2 < 0 \text{ set } X_2 = 0, \\ &\text{if } X_1 < 0 \text{ set } X_1 = 0, \\ &\text{if } X_1 + X_2 < N_2 \text{ set } X_2 = N_2 - X_1. \end{aligned} \tag{9}$$

The numbers of jailed and active citizens take only integer values, so X_1 and X_2 are floored to their nearest integers after each temporal step in the stochastic simulation of Eq. (3).

Figure 9 shows the time histories of X_1 and X_2 simulated from the effective reduced SDE. They qualitatively resemble closely Fig. 3 and clearly exhibit punctuated equilibria. We also plot the distribution of the outburst times (the times between two successive outbursts) and compare it to the one obtained from the original ABM (Fig. 10). The first three moments of the two distributions agree reasonably well (Table III). The probability density of the numbers of jailed and active citizens can now be computed by running the reduced simulation up to very long times (say, 10^7 time steps); there is little difference between the two simulation methods (Fig. 11). These small differences do not affect the validity of the effective reduced SDE to replicate the punctuated equilibrium state and its statistics. Moreover, simulation of the SDE for 2×10^5 steps takes about 45 min, whereas the original ABM requires over 17 h for the same simulation length (a difference factor of more than 20).

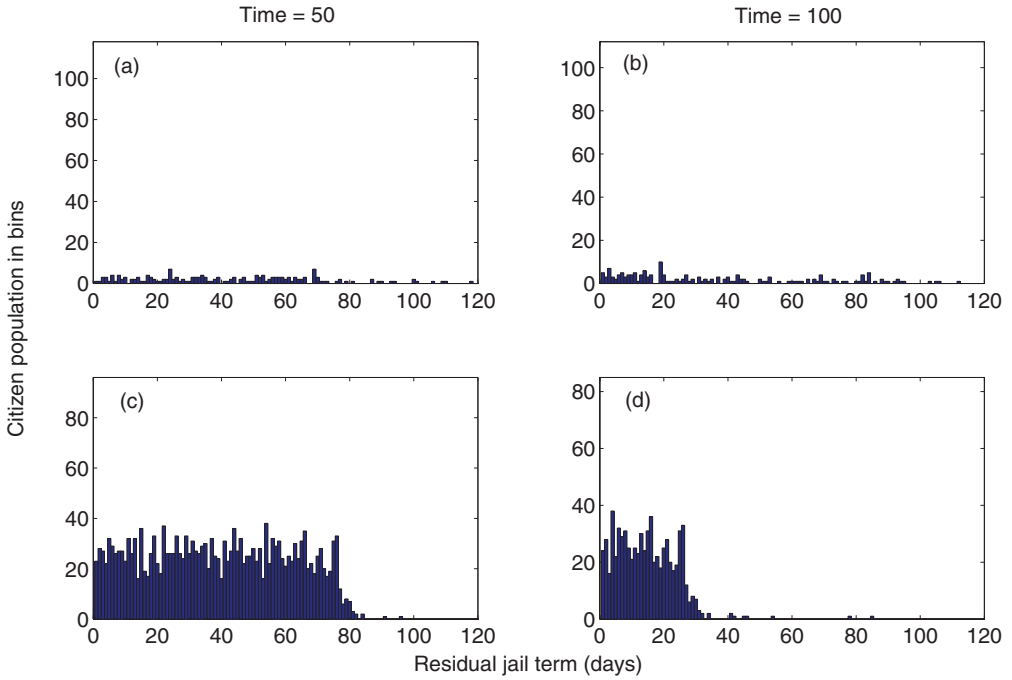

FIG. 8. (Color online) Distribution of residual jail terms for jailed citizens from regions 2 and 3 in Fig. 4(a): (a) distribution for region 2 jailed citizens at time 50, (b) distribution for region 2 jailed citizens at time 100, (c) distribution for region 3 jailed citizens at time 50, and (d) distribution for region 3 jailed citizens at time 100.

IV. JOINT PDF OF THE COARSE VARIABLES AND MEAN OUTBURST TIME: EFFECTIVE FOKKER-PLANCK AND BACKWARD KOLMOGOROV PARTIAL DIFFERENTIAL EQUATIONS

Now that an effective model in the form of a reduced SDE has been obtained and validated, a number of mathematical or computational tools, in the form of associated continuum partial differential equations, become available for the extraction of information about the system behavior and statistics.

The Fokker-Planck equation corresponding to the SDE (3) describes the evolution of the joint PDF of the numbers of jailed and active citizens $P(X_1,X_2,t)$:

$$\frac{\partial P}{\partial t} = -\sum_i \frac{\partial(\mu_i P)}{\partial X_i} + \frac{1}{2}\sum_{i,j} \frac{\partial^2(D_{ij}P)}{\partial X_i \partial X_j}. \tag{10}$$

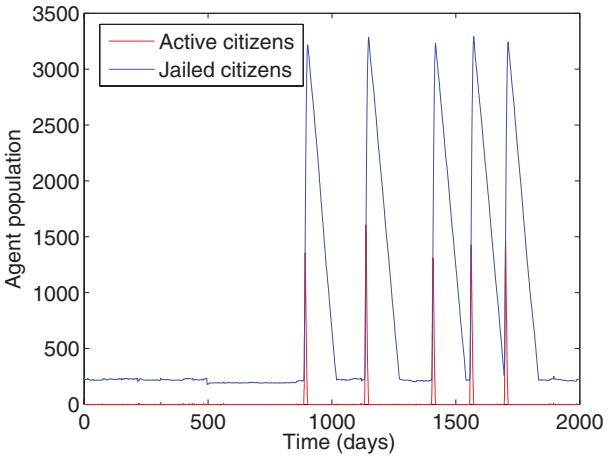

FIG. 9. (Color) Stochastic path of jailed and active citizens obtained from the simulation of the effective reduced SDE (3) subject to the boundary conditions (9).

Obtaining such a PDF directly from the SDE would require a large number of sample paths, that is, many stochastic realizations, each for time t, which can now be performed through a single, deterministic PDE computation. As $t \to \infty$, $P(X_1,X_2,t)$ approaches a stationary state [this can be visually verified by the fact that as the number of simulation time steps becomes larger, the joint PDF of (X_1,X_2) pairs appears very similar to that shown in Figs. 11(a), 11(c), and 11(e)]. At stationarity the left-hand side of the above PDE vanishes, leaving

$$-\sum_i \frac{\partial(\mu_i P)}{\partial X_i} + \frac{1}{2}\sum_{i,j} \frac{\partial^2(D_{ij}P)}{\partial X_i \partial X_j} = 0, \tag{11}$$

which we use to solve directly for the stationary distribution of the numbers of jailed and active citizens.

We note that no boundary condition corresponding to the boundary condition [Eq. (10)] of the SDE (3) is analytically known for the Fokker-Planck PDE. In order to circumvent this lack of explicit boundary conditions, we solve the Fokker-Planck equation in a much larger domain (approximately, the entire positive X_1 half plane). The values of the drift and diffusion coefficients in this extended domain come from the following considerations. (a) We set the drift coefficient μ_2 to positive values in the region below the simulation domain, as shown in Fig. 5; this would correspond to (X_1,X_2) trajectories of the SDE instantaneously bouncing back to

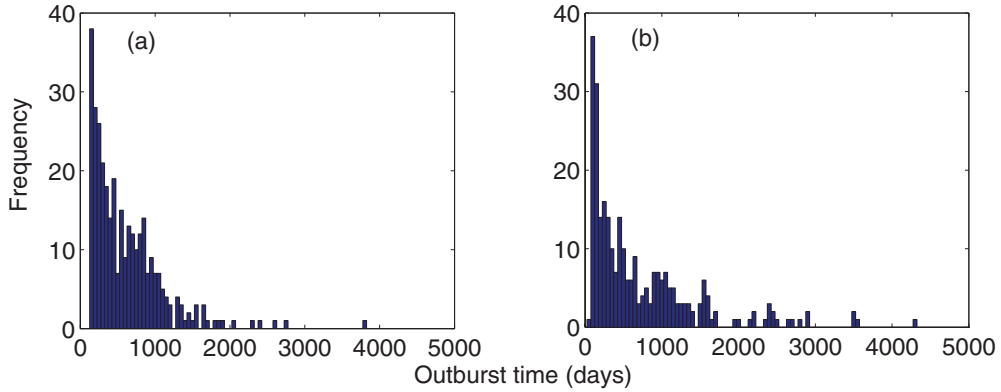

FIG. 10. (Color online) Comparison of the outburst time distributions from (a) the effective reduced SDE and (b) the original ABM.

the original domain if they attempted to exit it. (b) The diffusion coefficients D_{ij} in the same region have to be set appropriately to avoid significant spurious oscillations when Eq. (11) is solved. A large domain $\Omega = [0,300] \times [-100,100]$ encloses most observed trajectories. We set homogeneous von Neumann boundary conditions for the Fokker-Planck PDE on the boundary $\partial\Omega$. The PDE now reads

$$\begin{aligned} -\mathrm{div}(\mathcal{A}\nabla P + P\mathbf{b}) &= 0 \quad \text{in } \Omega, \\ (\mathcal{A}\nabla P + P\mathbf{b})\cdot\mathbf{n} &= 0 \quad \text{in } \partial\Omega, \end{aligned} \tag{12}$$

where ∇ represents the gradient and $\mathbf{n}$ stands for the outward unit normal vector to $\partial\Omega$. The matrix $\mathcal{A}$ and the vector $\mathbf{b}$ are given by

$$\mathcal{A} = \begin{pmatrix} -D_{11}/2 & -D_{12}/2 \\ -D_{12}/2 & -D_{22}/2 \end{pmatrix}, \tag{13}$$

$$\mathbf{b} = \left(\mu_1 - \frac{1}{2}\frac{\partial D_{11}}{\partial X_1} - \frac{1}{2}\frac{\partial D_{12}}{\partial X_2}, \mu_2 - \frac{1}{2}\frac{\partial D_{22}}{\partial X_2} - \frac{1}{2}\frac{\partial D_{12}}{\partial X_1}\right)^T.$$

The positive values for the drift coefficients in the region below the original domain are set to

$$\begin{aligned} &\mu_1 = 0, \quad \mu_2 = -X_2 \quad \text{if } X_1 > N_2,\ X_2 < 0; \\ &\mu_1 = 0, \quad \mu_2 = N_2 - X_2 \quad \text{if } 0 \leqslant X_1 < N_2,\ X_1 + X_2 < N_2. \end{aligned} \tag{14}$$

The diffusion coefficients in this region should have values comparable to those of the drift coefficients in order to avoid spurious oscillations. We adopt the idea of introducing additional isotropic artificial diffusion terms [19] to the left-hand side of Eq. (12) in the region $\Gamma = \{(X_1,X_2) : X_1 > N_2, X_2 < 0 \text{ or } 0 \leqslant X_1 < N_2, X_1 + X_2 < N_2\}$, i.e.,

$$\begin{aligned} D_{11} = D_{22} = -X_2/2, \quad D_{12} = 0 \\ \text{if } X_1 > N_2,\ X_2 < 0; \end{aligned}$$

$$\begin{aligned} D_{11} = D_{22} = (N_2 - X_2)/2, \quad D_{12} = 0 \\ \text{if } 0 \leqslant X_1 < N_2,\ X_1 + X_2 < N_2. \end{aligned} \tag{15}$$

We use the same finite-element formulation as in Ref. [20] to solve Eq. (12). Triangular elements on a uniform mesh are used for the discretization [Fig. 12(a)]. The resulting stationary $P(X_1,X_2)$ is normalized so that $\int_\Omega P(X_1,X_2)dX_1dX_2 = 1$. Figure 12(b) shows the PDF computed from Eq. (12), whose profile closely approximates that obtained from long SDE simulations [Fig. 11(e)].

Another continuum PDE whose solution provides important information for the behavior of stochastic solutions of our effective reduced SDE is the corresponding backward Kolmogorov equation. This equation can be used to obtain the mean exit time $\tau(X_1,X_2)$, which is defined as the mean time a path starting at (X_1,X_2) will take to exit a domain D in X_1-X_2 space; it reads (see, e.g., Refs. [20,21])

$$\sum_i \mu_i \frac{\partial \tau}{\partial X_i} + \frac{1}{2}\sum_{i,j} D_{ij}\frac{\partial^2 \tau}{\partial X_i \partial X_j} = -1. \tag{16}$$

The domain D of interest here is $[0,\infty] \times [-\infty,500]$. To avoid numerical difficulties in solving the very large set of linear equations resulting from the finite-element discretization of this entire domain, we confine the computational domain for Eq. (16) as $\tilde{D} = [0,600] \times [-50,500]$. A Dirichlet boundary condition is prescribed at $X_2 = 500$ and homogeneous von Neumann boundary conditions are used on the other sides of $\tilde{D}$. The reason we can set homogeneous von Neumann boundary condition at $X_2 = -50$ is because X_1 and X_2 in the negative X_2 domain can instantaneously bounce back to the simulation domain (Fig. 5), so the mean exit time for negative X_2 will be nearly uniform along the X_2 direction. The reformulated backward Kolmogorov equation is then given by

$$\begin{aligned} -\mathrm{div}(\mathcal{A}\nabla\tau) + \mathbf{b}\cdot\nabla\tau &= -1 \quad \text{in } \tilde{D}, \\ \tau &= 0 \quad \text{at } X_2 = 500, \\ (\mathcal{A}\nabla\tau)\cdot\mathbf{n} &= 0 \quad \text{at } X_1 = 0,\ X_1 = 600,\ X_2 = -50. \end{aligned} \tag{17}$$

We again use the finite-element formulation described in Ref. [20] to solve Eq. (17). The drift and diffusion coefficients at X_1 close to N_2 and X_2 close to zero have important contributions to the mean exit time $\tau(X_1,X_2)$ for all (X_1,X_2) pairs. We thus choose to use a nonuniform mesh, which has a (finer) size of 1 for $170 \leqslant X_1 \leqslant 180$ and $0 \leqslant X_2 \leqslant 5$ and

TABLE III. Comparison of the statistics (the first three moments) of the outburst time distribution between the true (ABM) and reduced (SDE) models.

Method	Mean	Standard deviation	Skewness
SDE	611.3	478.1	2.171
ABM	714.4	723.7	1.893

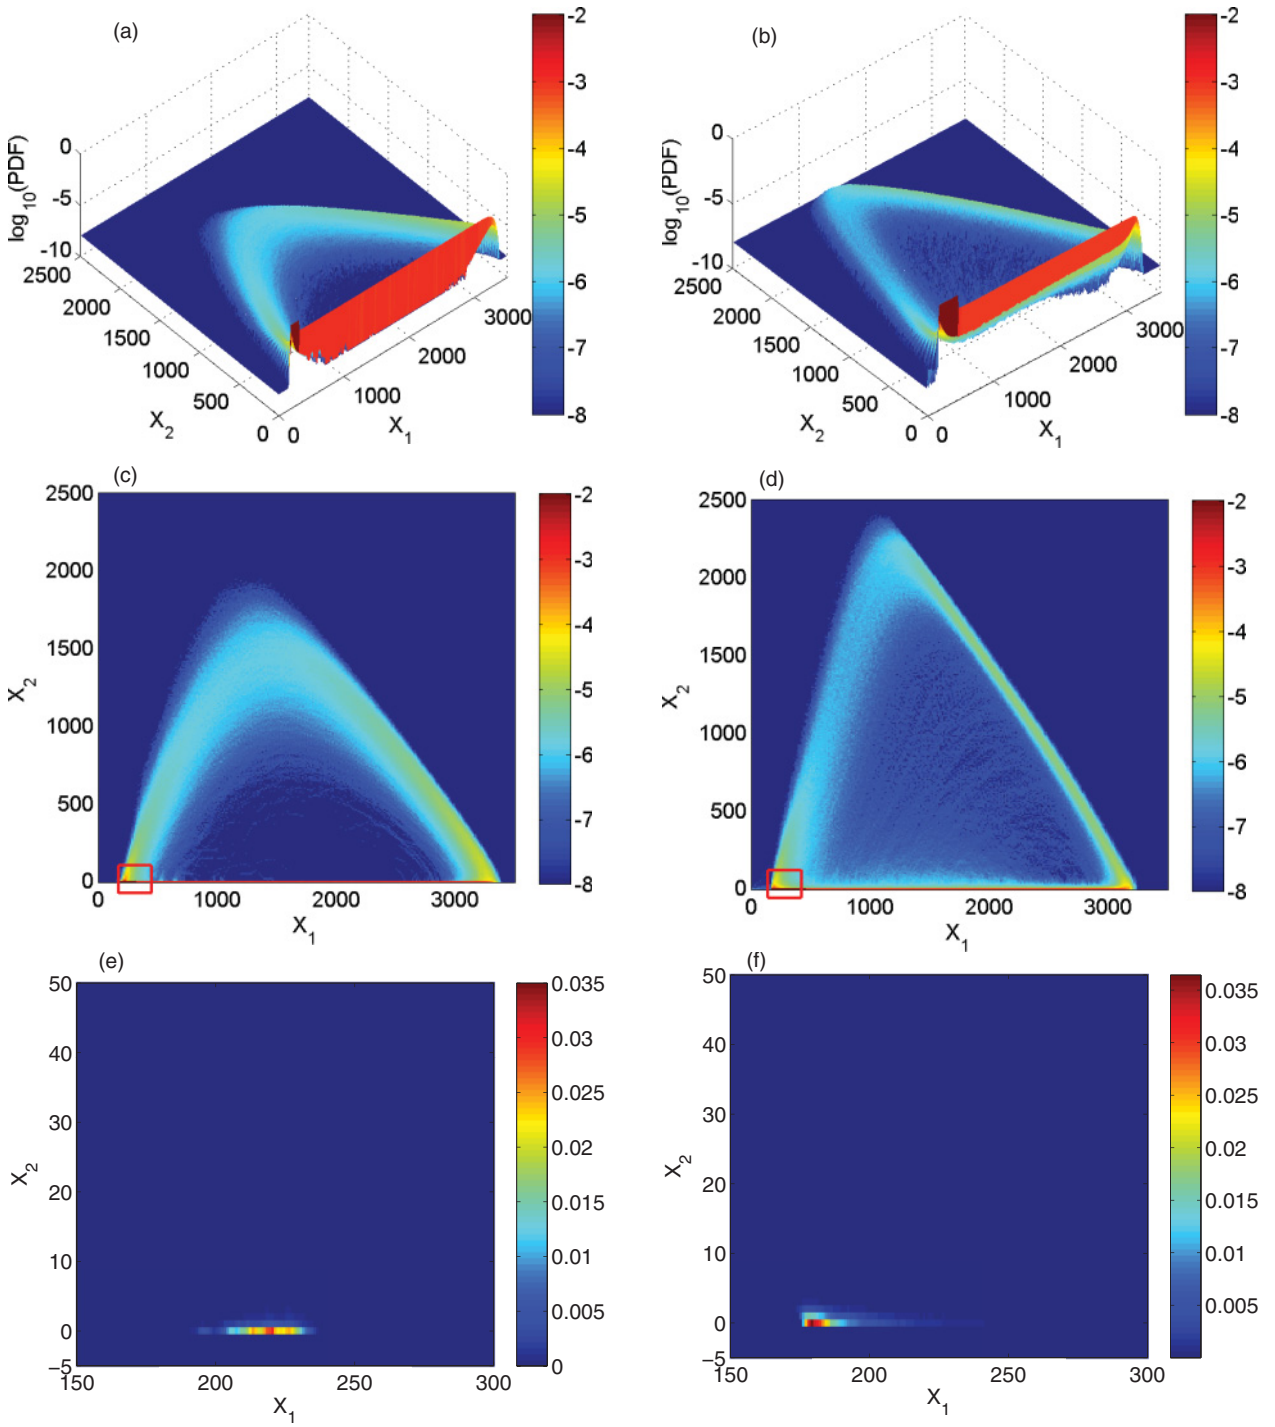

FIG. 11. (Color) Comparison of the PDF of the numbers of jailed and active citizens between the true (ABM) and the reduced (SDE) models. (a) Three-dimensional (3D) surface of $\log_{10}$ (PDF) simulated using the SDE, (b) 3D surface of $\log_{10}$ (PDF) simulated using the ABM, (c) 2D projected image of $\log_{10}$ (PDF) simulated using the SDE [the region in the red box is zoomed in and shown in (e)], (d) 2D projected image of $\log_{10}$ (PDF) simulated using the ABM [the region in the red box is zoomed in and shown in (f)], (e) zoomed-in image of the PDF simulated using the SDE in the subdomain $[150,300] \times [-5,50]$, and (f) zoomed-in image of the PDF simulated using the ABM in the same subdomain.

a (coarser) size of 5 elsewhere, to discretize the domain $\tilde{D}$. Triangular elements are used to solve Eq. (17) on this mesh [Fig. 13(a)]. The result for $\tau(X_1,X_2)$ is shown in Fig. 13(b). The mean exit time starting from $(X_1,X_2) = (600,0)$ is approximately 510. When running many sample paths of the effective reduced SDE we observe that the value of X_1 in the resulting time histories peaks at around 3250. The drift coefficient μ_1 at $X_2 = 0$ is approximately -25 for $600 \leqslant X_1 \leqslant 3250$; this would suggest that the trajectory of X_1 at such large X_1 values will mostly tend to go back toward

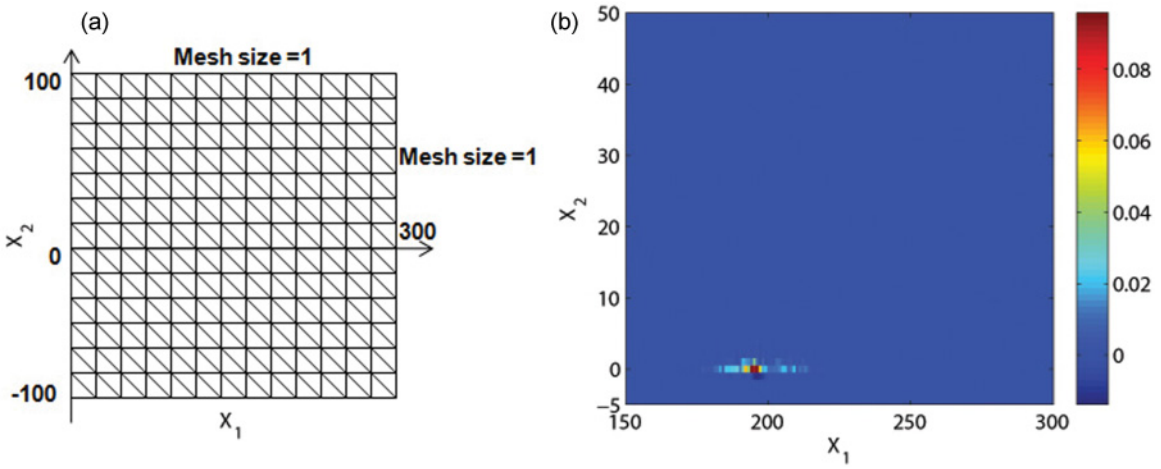

FIG. 12. (Color) (a) Finite-element mesh for solving Eq. (12) and (b) resulting stationary PDF of the numbers of jailed and active citizens.

$X_1 = 600$. Observing that an outburst (the rise to 3250) occurs almost instantaneously, using the above numbers to estimate the return to quiescence time by $(3250 - 600)/25 = 106$, and using the obtained value of 510 for the mean exit time, we obtain an estimate of $510 + (3250 - 600)/25 = 616$ for the mean outburst time, which agrees very well with the mean outburst time obtained from the SDE-based distribution of outburst times in Table III.

V. CONCLUSION

This paper explored the important issue of model reduction as a tool for the effective extraction of information from detailed, fine-scale agent-based models of social phenomena. In particular, using an intricate civil violence ABM as our illustrative example, we showed how to obtain an effective reduced model, here in the form of a two-degree-of-freedom SDE. An important consideration in such attempts at model reduction is the selection of the right variables: the (hopefully) small set of observables from the ABM simulation in terms of which a meaningful reduced description can in principle be obtained. In our example we found that two such coarse variables, the number of jailed citizens, and the number of active citizens could be used to formulate an effective reduced model.

Through an appropriately designed lifting procedure, which links variables between the reduced SDE (coarse) and the original ABM (fine) level of description of our complex system, and using the resulting coarse time stepper, we were able to estimate the drift and diffusion coefficients in the reduced effective SDE as functions of the two coarse variables only. The sample paths of the numbers of jailed and active citizens, as well as the statistics of the outburst time distribution from the reduced SDE model, closely resemble those obtained from the original (but significantly more time-consuming) agent-based simulation. We also demonstrated how to use continuum tools associated with the reduced SDE model, in particular the corresponding Fokker-Planck and backward Kolmogorov equations, to directly and efficiently extract additional useful information for the model behavior: We were able to directly compute the stationary PDF of the numbers of jailed and active citizens as well as to estimate the mean time between outbursts through the computation of mean exit times.

Equation-free techniques are intended to accelerate the (computational) extraction of information from a fine-scale

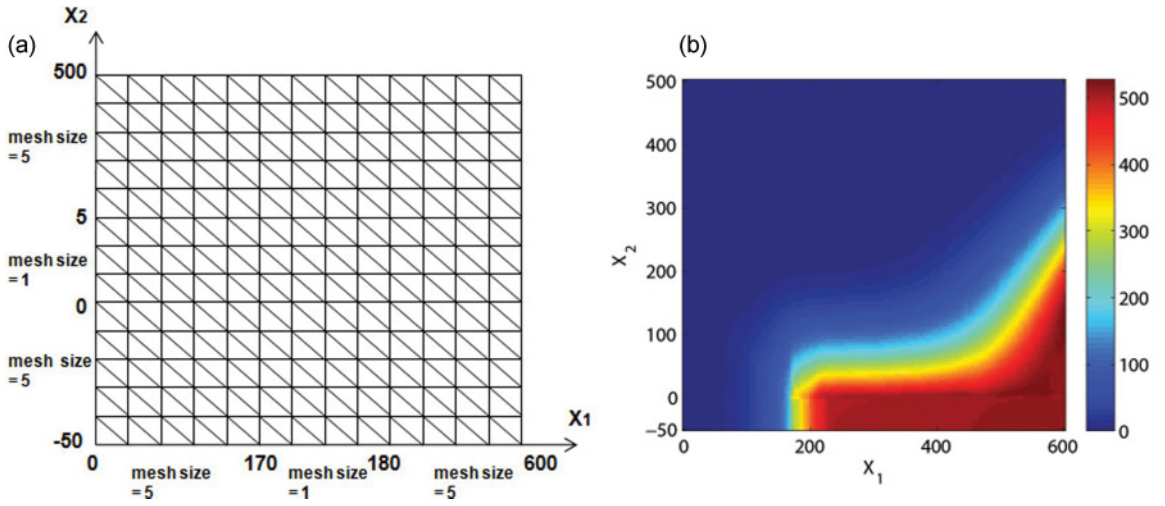

FIG. 13. (Color) (a) Finite-element mesh for solving Eq. (17) and (b) mean exit times computed as a solution of Eq. (17).

model (here an ABM interacting particle model). Coarse projective integration, for example, exploits the smoothness of the coarse behavior in time to compute with the ABM for only short time intervals; techniques such as patch dynamics exploit the smoothness of the coarse observables in space in order to compute in only parts of a model domain. To accurately gauge these savings one should also include the cost of repeated lifting and restriction; to be fair, one should also take into account the initial outlay of effort to determine what good coarse variables might be. It should be said, however, that certain tasks such as the location of unstable stationary states, or transition points (saddles), would be essentially impossible by direct simulation of the full ABM, so the amount of computational savings is clearly both problem dependent and, for a given problem, computational task dependent. The conceptual benefits from understanding that a complex model with many degrees of freedom can in principle be coarse grained are many and obvious: The equation-free approach is trying to capitalize on these benefits even when the coarse equations we postulate exist are not explicitly available.

Let us reiterate that qualitative observations from long-detailed simulations were incorporated in our lifting operator: Obtaining good fine-scale initial conditions for so many agents and states based on only two scalar values is a formidable task even when we suspect that only two coarse variables will suffice. It is important to note that the value of such a model reduction techniques lies only in the computational efficiency of extracting information from the ABM model; the correctness of the extracted information is completely determined by this original detailed ABM model. If the ABM is physically accurate, then the reduction approach can speed up the computational extraction of useful behavioral predictions and statistics. If, in contrast, the ABM model is wrong, then its (wrong) macroscopic consequences can be quickly found and thus indirectly help modify the model. It is important to observe that many of these reduction steps take the form of a wrapper: an algorithm that can be wrapped around the best agent-based model one has at hand. Changing the ABM, i.e., making it more detailed and/or more accurate, does not affect these wrappers, so they can be used without change. In our case the wrappers were the solvers of the two-dimensional SDE and the Fokker-Planck equation, with the values of the drift and diffusion coefficients in the equations provided by short agent-based simulation and estimation. For different problems and different types of tasks, these wrappers are motivated by different types of coarse-grained models and algorithms. If the coarse model is a deterministic ODE or PDE, initial-value solvers (such as the forward Euler algorithm) become the wrappers that help accelerate the temporal simulation of the unavailable coarse model; fixed-point solvers (based on matrix-free linear algebra such as Newton-Krylov generalized minimal residual method) become the wrappers that help locate coarse stationary states; and eigensolvers (such as Arnoldi-type algorithms) become the wrappers that help quantify the linearized stability of these states [22].

It is also important to note that in our illustrative example we already had (from experience with the simulations) a reasonably good idea of what good coarse variables might be. For problems where such *a priori* knowledge of the right macroscopic observables for the reduction is not available we expect that modern data-mining (manifold learning) techniques such as DIFFUSION MAPS [23,24] or ISOMAP [25] can be used to suggest such good reduction coordinates. It is worth noting that the structure of the equation-free approach (and in particular the lifting step) does not allow easily for coarse models with memory (e.g., rate-type models of neural activity). One must instead discover additional state variables (possibly higher-order spatial correlations) that embody the same information we would have in a model with fewer variables but also memory. Needless to say, it follows that the longer the memory, the more additional state variables we need (see the discussion in Ref. [26]).

We hope that the reduction approach we demonstrated here can be extended to the coarse graining of more detailed models that manifest the same punctuated equilibrium phenomenon in the social sciences as well as in evolutionary biology.

ACKNOWLEDGMENTS

We thank Dr. Radek Erban and Dr. Tomas Vejchodsky for their advice on solving the Fokker-Planck PDE. The work of V.A.F., M.F., and I.M. was partially supported by Air Force Office of Scientific Research (AFOSR) Grant No. FA9550-08-1-0217. Y.Z. and I.G.K. are grateful for AFOSR Grant No. FA9550-09-1-0018 and to the US Department of Energy for their support through Grant No. DE-SC0002097.

[1] J. von Neumann, *Theory of Self-Reproducing Automata* (University of Illinois Press, Champaign, 1966).
[2] T. Schelling, Am. Econ. Rev. **59**, 488 (1969).
[3] T. Schelling, Public Interest **25**, 61 (1971).
[4] T. Schelling, J. Math. Sociol. **1**, 143 (1971).
[5] K. Carley, Am. Sociol. Rev. **56**, 331 (1991).
[6] N. Gilbert and J. Doran, *Simulating Societies: The Computer Simulation of Social Phenomena* (University College London Press, London, 1994).
[7] R. Axelrod, *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration* (Princeton University Press, Princeton, 1997).
[8] R. Sun, *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation* (Cambridge University Press, Cambridge, 2006)
[9] K. Theodoropoulos, Y. Qian, and I. Kevrekidis, Proc. Natl. Acad. Sci. USA **97**, 9840 (2000).
[10] I. Kevrekidis, C. Gear, J. Hyman, P. Kevrekidis, O. Runborg, and C. Theodoropoulos, Commun. Math. Sci. **1**, 715 (2003).
[11] I. Kevrekidis, C. Gear, and G. Hummer, AIChE J. **50**, 1346 (2004).
[12] J. Epstein, Proc. Natl. Acad. Sci. USA **99**, 7243 (2002).
[13] B. Øksendal, *Stochastic Differential Equations: An Introduction with Applications* (Springer, Berlin, 2003).

[14] M. Fonoberova, V. Fonoberov, I. Mezic, J. Mezic, and P. Brantingham, J. Artif. Soc. Social Simul. **15**, 2 (2012).
[15] N. Eldredge and S. Gould, in *Models in Paleobiology*, edited by T. J. Schopf (Freeman, Cooper & Co., San Francisco, 1972), pp. 82–115.
[16] P. Bak and K. Sneppen, Phys. Rev. Lett. **71**, 4083 (1993).
[17] C. Laing, T. Frewen, and I. Kevrekidis, J. Comput. Neurosci. **28**, 459 (2010).
[18] Y. Aït-Sahalia, NBER Working Paper No. w8956 (unpublished).
[19] V. John and P. Knobloch, in *Proceedings of International Conference Programs and Algorithms of Numerical Mathematics 13*, edited by J. Chleboun, K. Segeth, and T. Vejchodsky (Academy of Science of the Czech Republic, Prague, 2006), pp. 122–136.
[20] R. Erban, S. Chapman, I. Kevrekidis, and T. Vejchodsky, SIAM J. Appl. Math. **70**, 984 (2009).
[21] R. F. Fox, Phys. Rev. A **33**, 467 (1986).
[22] C. I. Siettos, C. C. Pantelides, and I. G. Kevrekidis, Ind. Eng. Chem. Res. **42**, 6795 (2003).
[23] R. Coifman, S. Lafon, A. Lee, M. Maggioni, B. Nadler, F. Warner, and S. Zucker, Proc. Natl. Acad. Sci. USA **102**, 7426 (2005).
[24] R. Coifman and S. Lafon, Appl. Comput. Harmonic Anal. **21**, 5 (2006).
[25] J. Tenenbaum, V. de Silva, and J. Langford, Science **290**, 2319 (2000).
[26] C. Laing, T. Frewen, and I. Kevrekidis, Nonlinearity **20**, 2127 (2007).