Pre-Convergence Behavior and Quasi-Stationary Distributions: Framework and Applications*

Sebastian Bervoets† and Mathieu Faure‡

October 8, 2025

Abstract

We argue that looking at steady states in dynamical models often lacks relevance for economists when convergence time is excessively long, especially since steady-state outcomes and pre-convergence behavior can radically differ. To address this issue, we introduce the concept of *Quasi-Stationary Distributions*, which describe the behavior of a random process before convergence and we develop a method to characterize their support. We then illustrate our findings through three distinct models - on opinion formation, on online reviews, and on the persistence of false beliefs - emphasizing the stark contrast between long run and non-long run behavior.

Keywords: Quasi-stationary distributions, Non-long run behavior, Markov process

JEL Codes: C02, C65, D83

*We wish to thank Matthew Jackson, Daron Acemoglu, Fernando Vega-Redondo, Evan Sadler, Yves Zenou, Yann Bramoullé, Olivier Bochet, Hervé Crès, Alessandro Citanna for their useful comments, as well as participants of the several conferences where this work was presented. This work was supported by the French National Research Agency Grant ANR-17-EURE-0020, and by the Excellence Initiative of Aix-Marseille University - A*MIDEX. Mathieu Faure gratefully acknowledges the support of the French National Research Agency, under grant ANR CIGNE (ANR-15-CE38-0007-01).

†Corresponding author. Aix Marseille Univ, CNRS, EHESS, Centrale Marseille, AMSE, Marseille, France. Email address: sebastian.bervoets@univ-amu.fr. Postal address: Aix Marseille School of Economics - 5 Boulevard Maurice Bourdet - 13205 Marseille Cedex 01 - France.

‡Aix Marseille Univ, CNRS, EHESS, Centrale Marseille, AMSE, Marseille, France. Email address: mathieu.faure@univ-amu.fr

I Introduction

"The long run is a misleading guide to current affairs. In the long run we are all dead". This provocative quote from J.M. Keynes (Keynes (1923)) underlines two important features: first, the long run might materialize in such a long time that it makes little sense considering it; second, it suggests that the long run might look very different from the short or middle run, which are the time horizons that matter for making decisions.

In systems modeled as Markov processes with absorbing states, absorption time can be sufficiently long for the system to settle down before absorption, to what is called a quasi-stationary distribution (hereafter, QSD). In this paper, we introduce the concept of QSD and use it to develop tools to determine whether absorption time will be short or long. We also develop a method to characterize this QSD, which describes how the system behaves before it gets absorbed. We also show how the QSD and the absorbing states can be radically different, illustrating how it provides a better guide to current affairs than the long run.

To the best of our knowledge, the idea of analyzing pre-absorption behavior has not yet been introduced in economics.[1] This is surprising, since this idea emerged more than seventy years ago in mathematics (Yaglom (1947)), the concept of quasi-stationary distributions was formalized in the late fifties (Bartlett (1957), Bartlett (1960)) in the study of extinction in biological systems, and the theory was further developed in the early sixties (see Méléard et al. (2012) for an overview). Since then, QSDs have been used in various fields such as biology and ecology, chemistry, the study of epidemics, genetics, neuroscience... But they have not yet crossed the borders of social sciences, economics in particular. Our first contribution lies in introducing concepts from the QSD theory and defining a family of stochastic models - which can be used to describe many economic situations with local interactions - for which these concepts are relevant (Section III).

When the number of agents in a system is small, one can compute everything by hand and analyze the behavior of the system step by step. As the number of agents grows, the number of states within the system also expands, leading to the natural expectation that the absorption time correspondingly grows. Hence pre-absorption behavior becomes more and more important and QSDs might become more and more informative. Unfortunately, there is no general method to determine whether a QSD will in fact be relevant, nor to compute it in large systems. Our second,

[1] Some papers have already made the distinction between several time spans in the context of learning in games and selection of Nash equilibria. In Binmore and Samuelson (1994), followed by Binmore et al. (1995), the authors differentiate between *long run* and *ultra-long run* and argue that the selection process based on small random perturbations proposed by Kandori et al. (1993) and Young (1993) only works in the ultra-long run. They argue that any equilibrium can be understood as a long run outcome of the game, while the distribution of time spent in each equilibrium constitutes the ultra-long run. This is one main difference with what we do: we consider an non-ergodic system where the ultra-long run corresponds to absorbing states, and these outcomes are unrelated to long-run outcomes.

and main contribution, thus lies in identifying the conditions under which the QSD captures the process's pre-absorption behavior, and in approximating what it looks like.[2] In order to do that, we relate our random process to a deterministic dynamical system corresponding to the expected motion of the process when the size of the system is large. We then show (Theorems 1 and 2) that when this dynamical system admits at least one attracting state away from the absorbing states of the Markov process, then (i) absorption time will be exponential in the size of the system and the QSD will reflect pre-absorption behavior and (ii) the QSD of the system is concentrated around these attracting states which, are, in general, easy to determine.

Our third contribution is to illustrate these results by developing three applications. The diversity of these applications illustrates how QSDs can enlighten different topics, yet they share the common feature that the long run is easy to determine and offers predictions that seem intuitive. However, in the three cases we show that absorption time is too large to be realistic and we exhibit the QSD of the system, which features radically different properties from the long run, making the later deceptive for understanding social and economic behavior.

The first model we develop (section V) is a model of opinion formation, where agents influence each others through social interaction. Models of opinion formation offer a perfect framework to illustrate how QSDs are relevant. As stressed by Acemoglu and Ozdaglar (2011) in the introduction of their survey of opinion formation models: *"Like almost all of the literature in this area, [...] we will focus on the behavior of opinions and beliefs in a society in the "long run" [...]. It should be borne in mind throughout that the relevant "long run" might be longer than the lifespan of a single individual and might arrive so slowly that non-long run behavior might be of greater interest."* We agree and therefore focus on the non-long run of a simple model where agents from several communities interact and update their opinions over two candidates. This model will naturally end up with consensus, but we show that consensus is reached in a time that grows exponentially with the number of agents. Meanwhile, the system stabilizes for very long periods of time around a QSD before absorption into consensus. This QSD exhibits strong disagreement in society with roughly half of the population preferring one candidate and the other half preferring the other. It also involves a substantial polarization, with some communities strongly supporting one opinion and others strongly supporting the other. These features are radically different features from consensus, and are probably more realistic also.

The second model we develop (section VI) focuses on learning from online reviews, in the spirit of Acemoglu et al. (2022). Agents make purchase decisions for a given product based on both their ex-ante valuation of the item and the reviews left by previous customers. In their paper, Acemoglu

[2]As we state in section III, a QSD always exists. However, it may be both uninformative about the system's behavior and irrelevant, as the system might be absorbed before its distribution concentrates around it. It is therefore important to identify conditions under which the QSD is informative and relevant.

et al. (2022) provide conditions under which customers learn whether a product is of good or poor quality, a scenario referred to as “complete learning”. Complete learning occurs exponentially fast. In this context, we shift our attention to the scenario of incomplete learning, aiming to analyze whether good and bad products follow the same trajectory in the market. Our findings indicate that, once again, the asymptotic behavior is not a reliable predictor of non-long run behavior. Specifically, we show that the process governing the product rating admits a QSD whenever the product is of good quality, but not when it is of poor quality. Consequently, products of poor quality tend to disappear from the market almost immediately, while high-quality products persist for an extended period before eventually disappearing, all while maintaining a rating that closely reflects their quality.

In a third application (section VII), we explore the issue of persisting false beliefs or superstitions, such as the belief that the Earth is flat. Despite overwhelming evidence to the contrary, a small portion of the population still clings to this belief.[3] Without introducing distortions into natural models, causing certain agents to filter out specific types of information, it is difficult to explain why such false beliefs would not eventually disappear completely.[4]

Instead of introducing distortions, we follow another route by sticking to natural models in which all agents eventually acquire an accurate understanding of the true state of the world, and show that this learning process can be extraordinarily long. We then concentrate on pre-convergence behavior and exhibit the presence of a QSD, where a fraction of agents may persist in holding the belief that the Earth is flat, even though, in the long run, they will ultimately learn that it is in fact spherical.

The three applications highlight an advantage to focus on the QSD in random processes. It allows to explain real life observations while sticking to simple, natural models, without resorting to complex modifications of these models. This is the paper’s main insight, which echoes the conclusion of (Acemoglu and Ozdaglar, 2011)’s survey, as the authors say: *”The development of more powerful mathematical tools to study opinion dynamics away from the long run stationary distribution [...] constitutes an important area for future research.”*

We close the introduction with a discussion of related literature. A large strand of the literature in game theory has studied stochastic models of equilibrium selection, typically formulated as Markov chains with absorbing states perturbed by small amounts of noise. A seminal contribution is due to Kandori et al. (1993), who show that in large populations, the stochastically stable

[3]In fact, there exists a flat Earth society, which hosts conferences, publishes books, and maintains forums and blogs in an attempt to persuade skeptics that the Earth is not spherical. See https://theflatearthsociety.org.

[4]For instance, in Compte (2021), agents are assumed to have limited information-processing capabilities, thereby favoring information that confirms their initial (and potentially erroneous) beliefs. While more sophisticated agents will invariably come to understand that the Earth is spherical, less sophisticated ones may continue to believe it is flat.

equilibria - those selected in the long run under rare mutations - are precisely the ones that survive in the stationary distribution of the perturbed process. Similarly, Young (1993) develops the general framework of stochastic stability and adaptive learning, emphasizing how long-run play selects equilibria when agents occasionally make mistakes. In these models, the relevant set of candidate outcomes (the equilibria of the underlying game) is known ex ante, and the main question is which equilibrium will ultimately be selected. By contrast, our framework starts from a more general probabilistic object: a Markov chain with absorbing states, not necessarily derived from strategic interaction. In this case, the absorbing states are known but not where the process spends most of its time. Instead, the system lingers in regions that are not obvious a priori, and QSDs provide the appropriate tool to identify them.

Building on this, Ellison (1993) studies how the structure of interactions (global versus local) affects the speed of convergence to the stochastically stable equilibrium, while Kreindler and Young (2013) demonstrate that convergence can in fact be fast in large populations when the noise is moderate. Here again, our contribution departs from this literature by unraveling an unexpected form of metastability: rather than selecting among known equilibria, the QSD reveals regions of the state space where the process remains for exponentially long periods before absorption. This contrast also highlights a difference in temporal perspective: equilibrium selection focuses either on the ultra-long run, in which only the stochastically stable equilibria matter, or on the speed of convergence thereto. Our analysis instead emphasizes the pre-absorption "long run", which can persist for horizons that grow exponentially with the system size.

A related line of research explicitly distinguishes between the "long run" and the "ultra-long run." Binmore and Samuelson (1994), and later Binmore et al. (1995), argue that the long run corresponds to the stationary behavior induced by stochastic dynamics, while the ultra-long run describes the limiting distribution as the noise vanishes. In our framework, a similar distinction arises naturally: the "long run" is captured by the QSD, which describes the effective stationary distribution conditional on survival, while the "ultra-long run" corresponds to eventual absorption.

Taken together, this literature establishes a close connection between stochastic learning, equilibrium selection, and deterministic dynamics. We extend this connection by shifting the focus from equilibrium selection among known candidates to the identification of metastable distributions that are not known ex ante. In doing so, we provide a probabilistic characterization of the long run before absorption, which complements the classical view of the ultra-long run in equilibrium selection models.

II Motivating example

Consider a society consisting of two communities[5], each containing N agents. Each agent initially supports one of two candidates, A or B and opinions change with time: at each stage, one agent is taken at random. After meeting two other randomly drawn agents, he sticks to his initial opinion whenever he meets at least one other individual sharing his views, or he switches to the other candidate if both individuals he interacted with hold the opposite opinion. The two people he meets are drawn either from their own community, with high probability $(1 - p)$, or from the other community with low probability p. This fraction p measures the extent to which the two communities mix. We set $p = 0.1$, reflecting the fact that individuals interact more inside their own community than outside.

There are two absorbing states in this model, one in which everyone supports A, another one in which everyone supports B. Once this consensus is reached the system does not move anymore since no one will change opinion. The intuition of the dynamics are as follows: Assume the initial distribution of opinions is close to fifty-fifty in both communities. As the process of interactions starts, opinions should fluctuate around fifty-fifty in each community, until one community drifts away from fifty-fifty. This should then create a cascading effect within that community and a contagion effect to the other community, causing opinions to rapidly swing to everyone having the same opinion. The process just described could happen fairly quickly, since a drift can occur relatively soon and once it has started, the cascading effect should be fast.

However simulations suggest that these intuitions are not correct. First, absorption time seems to be exponential in N. We simulated a total of 1000 realization for each value of N and computed the average absorption time, represented in Figure 1. If it only takes a few thousand iterations to be absorbed into consensus when N is small, it takes on average more than 1.5 million iterations for a society with $N = 200$.

Second, when we look into pre-absorption behavior, we observe that the cascading effect does takes place, but it only takes place very abruptly at the end. Before that, the process does not seem to wander around various distributions of opinion, rather it seems to spend almost all its time around a fifty-fifty distribution. This is shown in Figure 2 where we see some realizations of the process, and absorption time growing exponentially with N.

These simulations actually show fluctuations around fifty-fifty situations, and the magnitude of these fluctuations is sometimes relatively large. For instance with $N = 50$, we observe several dips at around 35% of A voters, spikes at around 60%, and still the process goes back to fifty-fifty before eventually being absorbed. When N increases, these spikes and dips are less pronounced, but remain significant. This is surprising at first sight, as it would seem reasonable to expect

[5]The full-fledged version of this simplified model is developed in section V.

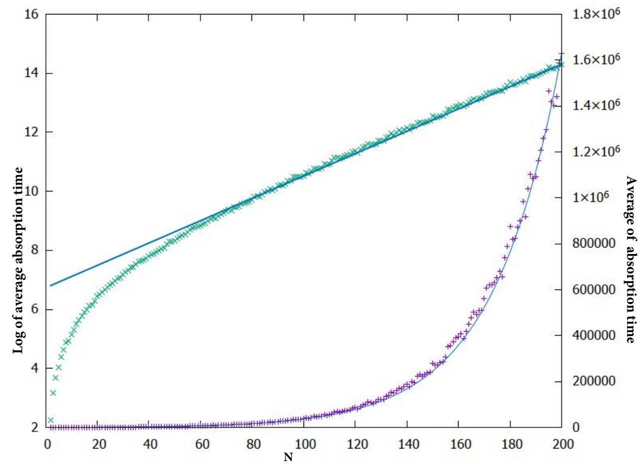

Figure 1: Absorption time as a function of N. The scale on the right shows average absorption time, corresponding to the purple dots, while the scale on the left shows the absorption time on a logarithmic scale, corresponding to the green dots. The blue lines are interpolations to fit the curves. The green dots follow a straight line when N is not too small.

that once the process drifts away from fifty-fifty, it is either absorbed rapidly, or fluctuates around another distribution of votes before continuing on a series of jumps, until absorption. But it seems that the process only goes back to fifty-fifty situations, as if this distribution played a special role.

Concentration around the fifty-fifty distribution does not result from inertia around the initial conditions. First, realizations which behave like those in Figure 2 can be reproduced even starting far from a fifty-fifty distribution, as illustrated in Figure 3. Second and as just described, even when starting close to fifty-fifty, there are quite large fluctuations early on, driving the process far from initial conditions.

The third observation is that the fifty-fifty situation only holds at the aggregate level. When we look into the behavior in each community, we see that neither community is evenly split between the two candidates. Instead, Figure 4 clearly shows that both communities have a large majority of either A or B supporters.

Actually, in each of these realizations the share of A voters remains around 85% in one community and 15% in the other before absorption[6]. This yields fifty-fifty on aggregate, but with a highly polarized society.

Taken together, our simulations indicate the following: that absorption time grows quickly with population size; that, conditional on non-absorption, the process leads to polarized societies; that, at any arbitrary point in time, the share of both types is close to fifty percent. Hence it

[6]When $N = 25$ and $N = 50$, polarization of society, i.e. the emergence of two communities very strongly favoring opposing candidates, is clearly shown by the graphs during the first few iterations. Note that, for N larger than 100, the transition from homogeneous communities to polarized communities cannot be observed on the graphs, because of the very large scales of the x-axis.

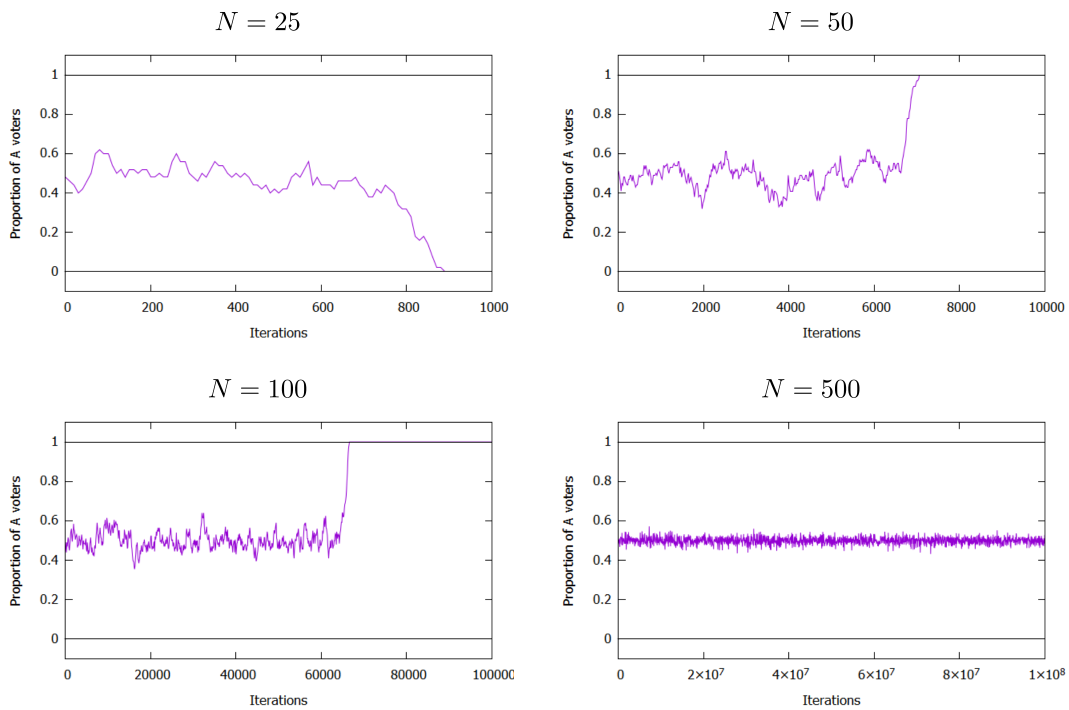

Figure 2: Aggregate proportion of A voters, for N varying from 25 to 500, when $p = 0.1$. The scale changes for each panel. Initial conditions are drawn at random between 0.4 and 0.6 in each community. Note that when $N = 500$, the process is still not absorbed after more than 10 billion iterations.

appears that focusing on absorbing states, which consist of unanimous societies, to predict society's behavior is not satisfactory. The relevant information is contained in the pre-absorption behavior and, in what follows, we provide a methodology to identify this pre-absorption behavior.

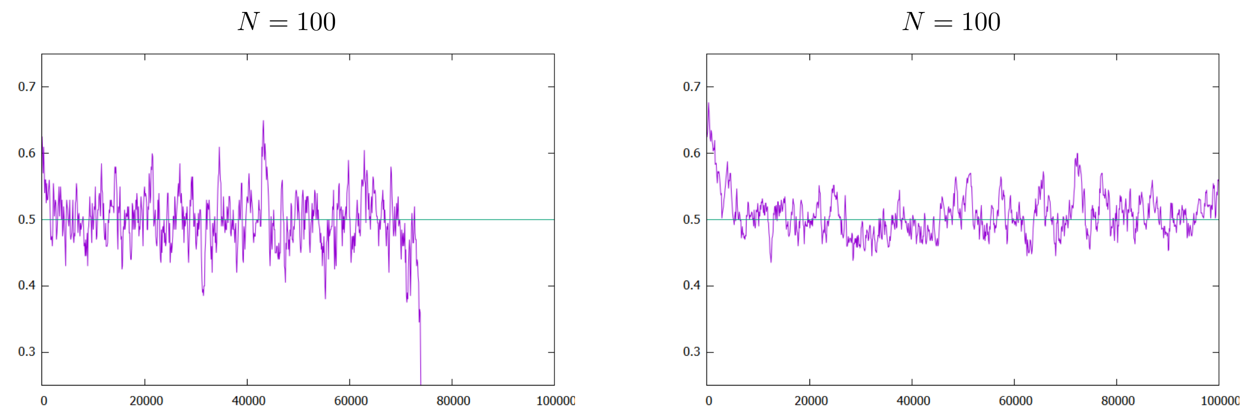

Figure 3: Two realizations with $N = 100$ and $p = 0.1$, with initial conditions of 40% of A voters in community 1 and 85% of A voters in community 2. The 50% line is represented in green.

III The model

Let $\mathbf{S}$ be a compact convex subset of $\mathbb{R}^k$. We consider a family of $\mathbf{S}$-valued random processes, indexed by $N \in \mathbb{N}^*$, which we denote by $\{\mathbf{x}^N\}_{N\geq 1}$. Typically, the dimension of the state space, k, represents the number of possible attributes, while N can be thought of as the "size" of the population. Hence $\mathbf{x}_n^N = (x_{n,1}^N, ..., x_{n,k}^N)$ can be thought of as the vector of frequencies of the different types in a population of size N, at period n.

Given $N \in \mathbb{N}^*$, we assume that the random process $(\mathbf{x}_n^N)_{n\geq 0}$ takes values in $\mathbf{S}^N := \mathbf{S} \cap \left(\frac{1}{N}\mathbb{Z}^k\right)$, a discrete state space, and can be recursively written as a stochastic difference equation of the form

$$\mathbf{x}_{n+1}^N - \mathbf{x}_n^N = \frac{1}{N} Z(\mathbf{x}_n^N), \tag{1}$$

where, given $\mathbf{x} \in \mathbf{S}$, $Z(\mathbf{x})$ is a random variable taking values in $\mathbf{D} := \{-1, 0, 1\}^k$, corresponding to the possible jumps of the process.[7] We denote by $\mu(\cdot \mid \mathbf{x})$ the distribution of $Z(\mathbf{x})$, $\mathbf{S}_0 := \{\mathbf{x} \in \mathbf{S} :\ \mu(\mathbf{0} \mid \mathbf{x}) = 1\}$ the set of absorbing states,[8] and $\mathbf{S}_* := \mathbf{S} \setminus \mathbf{S}_0$ the set of transient states.

Assumption 1

- *For all* $N \in \mathbb{N}^*$*, the set* $\mathbf{S}_0^N := \mathbf{S}_0 \cap \left(\frac{1}{N}\mathbb{Z}^k\right)$ *is nonempty.*
- *For all* $\mathbf{x} \in \mathbf{S}_*$*, Supp* $(\mu(\cdot \mid \mathbf{x})) = \{\mathbf{d} \in \mathbf{D} : \exists \epsilon > 0$ *such that* $\mathbf{x} + \epsilon\mathbf{d} \in \mathbf{S}\}$.
- *Given* $\mathbf{d} \in \mathbf{D}$*, the map* $\mathbf{x} \mapsto \mu(\mathbf{d} \mid \mathbf{x})$ *is continuous on* $\mathbf{S}_*$.

[7] We could use $\mathbf{D} := \prod_{i=1}^k \{-a_i, ..., -1, 0, 1, ..., b_i\}$ for some $a_i, b_i \in \mathbb{N}^*$ instead, for a more general model without changing any of the results. We restrict to the simplest case here, which is sufficient for our applications.

[8] Instead of considering absorbing states, we could consider the set of recurrent classes without changing our results. We use absorbing states, i.e. classes of size one, for simplicity.

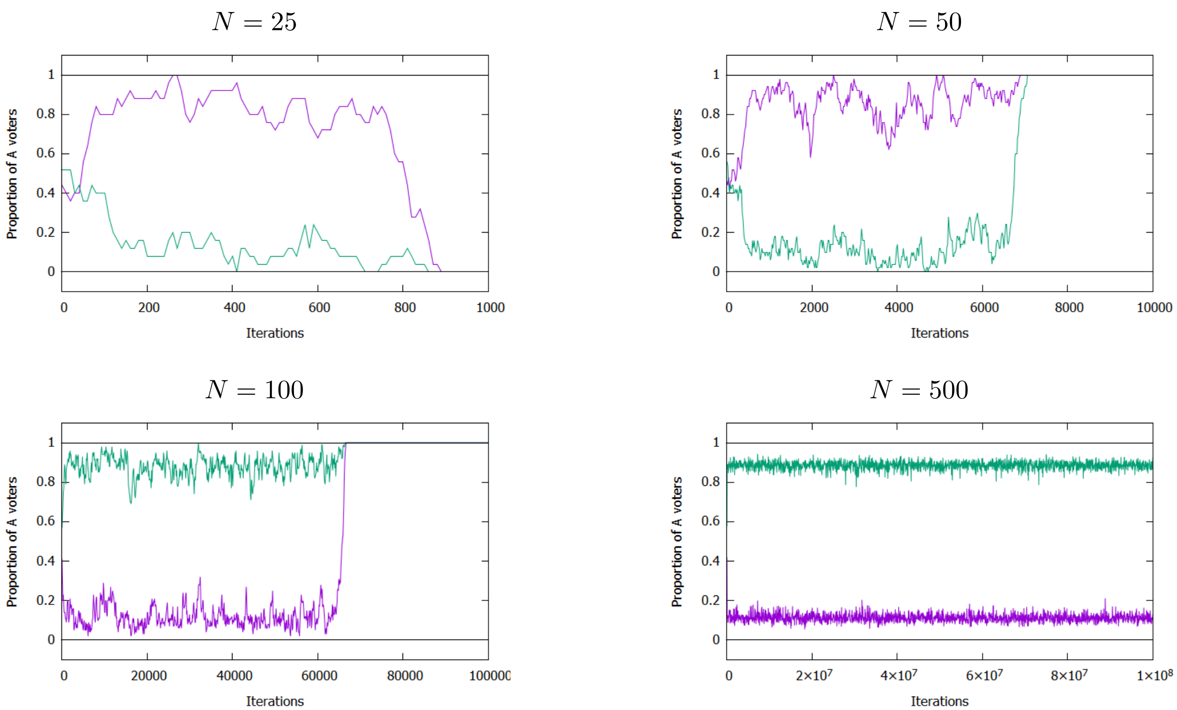

Figure 4: Proportion of A voters in each of the two communities. Each panel corresponds to the realization shown in Figure 2.

The first point says that there are absorbing states for each N, and at least some are common to every N, eventhough the lattice on which $\mathbf{x}^N$ evolves depends on N. Note that $\mathbf{S}_0$ is not necessarily finite, even though $\mathbf{S}_0^N$ is for all N. The second point assumes that from any transient state, the system can always go in every *feasible* direction (including remaining at the same position). This implies that $\mathbf{Q}_N$, the transition kernel associated to $\mathbf{S}_*^N$, is irreducible and aperiodic and therefore, given any N, the Markov process $\mathbf{x}^N$ will get absorbed in finite time T^N almost surely:

$$\forall \mathbf{x} \in \mathbf{S}^N, \ \mathbb{P}_{\mathbf{x}}(T^N < +\infty) = 1, \ \text{ where } \ T^N := \min\{n > 0 : \ \mathbf{x}_n^N \in \mathbf{S}_0^N\}. \tag{2}$$

Despite unavoidable absorption, the process might spend a long time in the transient space as N gets large. *Quasi-stationary distributions* (hereafter, QSD) provide insights on the *meta-stable behavior* before absorption, i.e. the potential stabilization of the process in the non-long run. A QSD is an invariant distribution, conditional on the process not being absorbed yet:

Definition 1 *Let $(\mathbf{x}_n)_n$ be a random process on $\mathbf{S}$. A probability measure π on $\mathbf{S}_*$ is a QSD for $(\mathbf{x}_n)_n$ if, for any $\Gamma \subseteq \mathbf{S}_*$ and any n,*

$$\mathbb{P}_\pi \left(\mathbf{x}_n \in \Gamma \mid \ \mathbf{x}_n \in \mathbf{S}_* \right) = \pi(\Gamma), \tag{3}$$

where $\mathbb{P}_\pi(\cdot)$ is the probability conditional on $\mathbf{x}_0$ being distributed according to π.

The theory of QSDs stems from the 60's. Here we present standard results from Darroch and Seneta (1965) for finite state Markov chains, adapted to our framework.[9]

First, the Markov process $\mathbf{x}^N$ admits a *unique* QSD for any $N \in \mathbb{N}^*$, since $\mathbf{Q}_N$, the restriction of the transition matrix to the transient space is an irreducible aperiodic kernel on the finite space $\mathbf{S}^N$. We call this QSD, π_N. Second, π_N has full support on $\mathbf{S}^N_*$, and is such that

$$\pi_N \mathbf{Q}_N = \lambda_N \, \pi_N \tag{4}$$

where $\lambda_N := \mathbb{P}_{\pi_N}(T^N > 1)$, is the probability that the process is not absorbed in one step, conditional on being initially distributed as π_N. Third, if initial conditions are distributed as π_N, then T^N is geometrically distributed, and

$$\mathbb{E}_{\pi_N}[T^N] = \frac{1}{log(\frac{1}{\lambda_N})}. \tag{5}$$

Fourth, we have the ergodic property: given $N \in \mathbb{N}^*$,

$$\lim_{n \to +\infty} \mathbb{P}_{\mathbf{x}} \left(\mathbf{x}^N_n \in \Gamma \mid T^N > n \right) = \pi^N(\Gamma), \;\; \forall \mathbf{x} \in \mathbf{E}_*, \; \forall \Gamma \subseteq \mathbf{E}_*. \tag{6}$$

As we see, the QSD π_N is the left eigenvector associated to $\mathbf{Q}_N$, with eigenvalue λ_N. In our context, these results convey two principal messages. First, the absorption time depends on how large λ_N is, where λ_N can be interpreted as the probability that the process does not get absorbed after one step. If λ_N is "much smaller"[10] than one, the process will be absorbed in a few steps, and therefore analyzing its QSD is irrelevant. If λ_N is "very close" to one, the absorption time will be "extremely long", and analyzing its QSD becomes critical to understand the economic phenomena under scrutiny.

Second, note that according to (3), a process distributed as π_N will remain distributed as π_N, as long as it remains in the transient space. However, if the process was never distributed as π_N in the first place, the non-absorbed process could look totally different from the QSD. The ergodic property rules out that possibility since pre-absorption behavior does not depend on initial conditions, as soon as absorption time is "long".

IV Absorption Time and Pre-Convergence Behavior

The random process $(\mathbf{x}^N_n)_{n \geq 0}$ defined in (1) can be recursively written as the following constant step-size stochastic difference equation

$$\mathbf{x}^N_{n+1} - \mathbf{x}^N_n = \frac{1}{N} \left(f(\mathbf{x}^N_n) + U^N_{n+1} \right) \tag{7}$$

[9] A formal statement of these claims is provided in Section IX.1, and in particular in Theorem A. The interested reader willing to go further will find a thorough exposition of the theory of QSDs in Méléard et al. (2012).

[10] All terms between quotes in this paragraph will be made precise in the next sections.

where $f : \mathbf{S} \to \mathbf{S}$, the drift, is given by

$$\mathbf{x} \mapsto \sum_{\mathbf{d} \in Supp(\mu(\cdot|\mathbf{x}))} \mu\left(\mathbf{d} \mid \mathbf{x}\right) \mathbf{d}$$

and U^N is a uniformly bounded martingale difference.

This recursive writing allows us to separate the drift, which is a deterministic term, from the random term, U^N, which is 0 on average. This guarantees that the random process $(\mathbf{x}_n^N)_{n \geq 0}$ is correctly approximated by the solutions of the deterministic system:

$$\dot{\mathbf{x}}(t) = f(\mathbf{x}(t)) \tag{8}$$

Let $(\phi(\mathbf{x}, t))_{t \geq 0, \mathbf{x} \in \mathbf{S}}$ denote the semi-flow of (8), and $\mathcal{E}$ denote the set of equilibrium points, i.e. zeroes of f. For simplicity, we assume that the set of equilibrium points is finite and that, for any initial condition $\mathbf{x} \in \mathbf{S}$, there exists some equilibrium $\hat{\mathbf{x}} \in \mathcal{E}$ such that $\lim_{t \to +\infty} \phi(\mathbf{x}, t) = \hat{\mathbf{x}}$.[11]

We say that $\hat{\mathbf{x}} \in \mathcal{E}$ is an *asymptotically stable* equilibrium for the flow ϕ if there exists an open neighborhood U of $\hat{\mathbf{x}}$ such that

$$\lim_{t \to +\infty} \sup_{\mathbf{x} \in U} \|\phi(\mathbf{x}, t) - \hat{\mathbf{x}}\| = 0.$$

The union of all open sets with this property is called the *basin of attraction* of $\hat{\mathbf{x}}$. Let $\mathcal{A} := \{\hat{\mathbf{x}} \in \mathcal{E} \cap \mathbf{S}_* : \hat{\mathbf{x}} \text{ is asymptotically stable}\}$ denote the set of asymptotically stable equilibria away from absorbing states.

We are now ready to state our two main results, relating the pre-absorption behavior of the random process to the set of equilibria of the system (8).

Theorem 1 *Assume that $\mathcal{A} \neq \emptyset$. Then there exists a constant $C_{\lambda_N} > 0$ such that $\lambda_N \geq 1 - e^{-C_{\lambda_N} N}$. We thus have*

$$\mathbb{E}_{\pi_N}(T^N) \geq e^{C_{\lambda_N} N}. \tag{9}$$

When the underlying deterministic system admits at least one asymptotically stable equilibrium, Theorem 1 implies that λ_N approaches one very fast as N grows. As a result, the expected absorption time increases at an exponential rate with the population size. The proof of Theorem 1 builds on the theory of random perturbations of dynamical systems, carefully adapted to our setting. Exiting a neighborhood of some $\hat{\mathbf{x}}$, despite the fact that the deterministic dynamics never leave it, has very small probability. Such an exit requires a large deviation from the deterministic flow, which can only occur through a sequence of unlikely unfavorable realizations of the

[11] This assumption can be relaxed, allowing for instance the existence of *attractors* (i.e. more general attracting sets, see the definition in the proof section in the appendix).

noise term U^N, which is a uniformly bounded martingale difference. Exponential martingale inequalities are then used to establish the exponential upper bound on this probability. Finally, since the QSD has full support, the process almost surely visits such a neighborhood, yielding the result.

By establishing an exponentially long absorption time, Theorem 1 justifies focusing on pre-convergence behavior, which we do now. For the pre-convergence behavior to be relevant, however, we must make sure that the process does not exhibit the peculiar behavior of spending a very long time in the vicinity of the absorbing space, without being absorbed. If that could happen, the process could wander around an absorbing state without being absorbed with very high probability, and the QSD might be artificially taken out of the picture. We discard this possibility with the following very mild assumption:

Assumption 2 $\forall \alpha < 1$, *there exists an open neighborhood V_α of $\mathbf{S}_0$ such that, for N large enough,*

$$\mathbb{P}\left(\mathbf{x}_N^N \in \mathbf{S}_0 \mid \mathbf{x}_0^N = \mathbf{x}\right) \geq \alpha^N, \ \ \forall \mathbf{x} \in V_\alpha.$$

Thus, once in a vicinity of an absorbing state, the process will get absorbed shortly with very high probability. If the process spends a large amount of time around some state before being absorbed, it is far from the absorbing states. It is therefore relevant for understanding pre-convergence behavior.

We are interested in the behavior as N increases. We thus consider the sequence $(\pi_N)_N$, where π_N is the unique QSD of the system of size N, and let N grow. The limit of this sequence is called a *limiting measure.*

Definition 2 (Limiting measure) *We call limiting measure any weak*-limit point*[12] *of the sequence $(\pi_N)_N$.*

We then prove the following:

Theorem 2 *Assume that $\mathcal{A} \neq \emptyset$, and let π^* be a limiting measure. Then Supp$(\pi^*) \subseteq \mathcal{A}$.*

This, together with Theorem 1, constitutes our main result. It says that in systems with many individuals, not only will absorption take an extremely long time, but the system will spend almost all that pre-absorption time around stable equilibria of the underlying dynamical system. We know a unique QSD exists, but computing it is not feasible. Our result provides a way to get round this problem and characterize its behavior.

[12] A probability measure $\pi^* \in \mathcal{P}(\mathbf{S})$ is a weak*-limit point of the sequence $(\pi_N)_N$ if there exists a subsequence $(\pi_{N_p})_p$ which converges to π^* for the weak* topology: $\lim_{p\to+\infty} \int f d\pi_{N_p} = \int f d\pi^*$ for any continuous function f on $\mathbf{S}$.

The main challenge in proving this theorem lies in the following observation. As shown in the proof of Theorem 1, whenever $\mathbf{x}_N^N$ lies in a neighborhood of an asymptotically stable equilibrium of the flow ϕ, the realizations of the process remain close to the solution curves of ϕ with high probability. This was established by deriving an exponential upper bound on the probability of leaving such a neighborhood. However, Theorem 2 shows that these are the only neighborhoods where this property holds. In particular, unstable equilibria are also points where the deterministic dynamics remain forever, yet the stochastic process must be able to escape their neighborhoods. We thus need to establish that in the vicinity of unstable equilibria it is sufficiently likely for the process to deviate from the flow and eventually exit the neighborhood. This is done by providing a lower bound on the probability that a sufficiently large deviation occurs.

IV.1 Comments on Theorems 1 and 2

Our main theoretical results call for important remarks to clarify what we claim and what we do not.

Convergence time. Three important points should be noted here. Firstly, Theorem 1 is about *expected* absorption time. This means that for some realizations, the process can be absorbed quickly. This might occur due to specific sequences of random draws or initial conditions close to an absorbing state. It will also happen if the process follows a trajectory that does not enter the neighborhood of any stable point or set. In fact, when following those trajectories, absorption time grows polynomially or even linearly with N, not exponentially. The exponential scaling of expected absorption time occurs only because the system can become trapped around points in $\mathcal{A}$, where it spends a significant amount of time before escaping.

Secondly, one could question the relevance of Theorem 2 if the stochastic process took an extensive amount of time to reach the neighborhood of a stable equilibrium, since we do not know how it would behave during that time. However, this will not happen, since the connection with the deterministic approximation ensures that the stochastic process follows the trajectories of the dynamical system. Thus Theorem 2 describes the behavior of the process after a relatively short time. In fact, stable equilibria and absorbing states play similar roles with respect to the process. The nature of these points, whether stable equilibria or absorbing states, does not impact the number of steps it takes for the random process to get close. The process will either go quickly in a small neighborhood of one or the other.

Thirdly, one might argue that the neighborhoods of absorbing states, rather than the absorbing states themselves, are the relevant objects. For example, in the opinion formation illustration, it makes little difference whether everyone holds the same opinion or if a large fraction (close to one) does. One could therefore define a relevant neighborhood around each absorbing state and consider the hitting time to this neighborhood instead of focusing on T^N. However, this distinction is largely irrelevant: Assumption 2 ensures that once the process enters the neighborhood of an absorbing state, absorption occurs rapidly. In fact, the phenomena captured by the QSD take place "far" from the set of absorbing states.

Characterization of the QSD. The ergodic property (6) suggests that - conditional on non-absorption after a very long time - the system ends up being distributed as its QSD, irrespective of initial conditions. This independence from initial conditions is called the mixing property. Yet, there is a possibility that the mixing time, the duration required for the mixing property to take effect, is significantly longer than the expected absorption time, in which case the QSD might not be relevant. What we ascertain is that, as long as absorption does not occur, the Markov chain can only spend time in regions where the QSD puts its weight. There are no other regions where it could remain for significant durations.

Also, our main result identifies the support of the QSD, yet it does not provide the exact distribution of this QSD. The stable equilibria that form the support of the QSD play in fact the same qualitative role as the absorbing states. Thus, in the same way that standard theory allows us to describe the set of absorbing states without being able to predict which one will be realized, what Theorem 2 provides is the set of states around which the system is likely to be trapped for a long time. However, each state within the support of the QSD has a positive probability of being visited by the process. Therefore, our characterization of this set is minimal in the sense of inclusion of sets.

Finally, although a QSD always exists, it may be irrelevant for understanding the system's behavior. This occurs when the associated dynamical system has no stable equilibrium. In such cases, absorption takes place rapidly, preventing the system from concentrating around the QSD; and even if it did, the QSD itself would not display any distinctive structure. This situation is illustrated in Section VIII. Hence, even when the QSD can be derived analytically or approximated numerically, our results remain insightful and reveal whether the QSD provides a meaningful description of the system or not.

V Application 1: A Model of Opinion Formation

In this application we use a model in which individuals always end up sharing a consensus in the very long run, while exhibiting polarization of opinions between communities for a very long time before. It is the model illustrated in section II, which generates persistent disagreement without the need to twist standard opinion formation models.

V.1 Description of the model

Consider a society consisting of 2 *communities* (or *groups*), each composed of $N \in \mathbb{N}$ individuals[13]. Every individual is either of type A or B (in terms of, for example, what candidate to support in an election, which of two different incompatible technologies to choose, whether or not to believe in a rumor...). Time is discrete and at period n, the state of the population is described by $\mathbf{x}_n^N = (x_{n,1}^N, x_{n,2}^N)$, where $x_{n,i}^N \in [0,1]$ is the fraction of type A individuals in community i.

At each period, one individual is drawn at random from each community to interact with other people, either from own community or from the other one. After interaction, the individual maintains or changes his opinion depending on the result of the interaction and a new period starts. We can think of many possible ways in which interactions take place, and/or in which decisions to maintain or change opinions are made. In this application, we assume that individuals interact with people from their own community with high probability $(1-p)$, and with people from the other community with lower probability, p, with $p \in (0,1)$. We also assume that individuals meet two other individuals and only change their opinions when these two individuals both share the opposite opinion[14]. Thus the probability that a B supporter switches to A is

$$p_i(\mathbf{x}) = (1-p)x_i^2 + px_{-i}^2 \tag{10}$$

Clearly, as long as there are individuals of opposite opinion in society there is always at least an individual who can change opinions with positive probability. Thus there are only two absorbing states, one where every individual supports A, and the other where every individual supports B.

V.2 Existence of a QSD and Deterministic Approximation

The random sequence $(\mathbf{x}_n^N)_{n\geq 0}$ is a Markov chain on $\mathbf{S^N} = \{0, \frac{1}{N}, \frac{2}{N}, ..., 1\}^2$ and $\mathbf{S_0} = \{\mathbf{0}, \mathbf{1}\}$. It can be recursively written as the random difference equation

$$\mathbf{x}_{n+1}^N - \mathbf{x}_n^N = \frac{1}{N} Z(\mathbf{x}_n^N) \tag{11}$$

where $Z(\mathbf{x}) = (Z_1(\mathbf{x}), Z_2(\mathbf{x}))$ is distributed as follows:

$$Z_i(\mathbf{x}) = \begin{cases} +1 & \text{with probability} \quad (1-x_i)p_i(\mathbf{x}), \\ -1 & \text{with probability} \quad x_i p_i(\mathbf{1}-\mathbf{x}), \\ 0 & \text{with probability} \quad 1-(1-x_i)p_i(\mathbf{x}) - x_i p_i(\mathbf{1}-\mathbf{x}) \end{cases}$$

[13]Many assumptions, such as the number of communities or communities having the same size, can be relaxed without difficulty. We describe the simplest settings for illustration purposes.

[14]In the appendix we provide a micro-foundation to this behavior, where individuals choose who to support by maximizing a suitable payoff function.

Let $\mu(\cdot \mid \mathbf{x})$ be the distribution of $Z(\mathbf{x})$. In this application, $\mathbf{D} = \{-1, 0, 1\}^2$, and $\mathbf{x} \mapsto \mu(\cdot \mid \mathbf{x})$ is continuous if p_i is, thus Assumption 1 is satisfied. We prove in the appendix (see Lemma 1), that Assumption 2 also holds.

The drift associated to the random process (11) is obtained by computing the expected value of $Z(\mathbf{x})$. In this case, in each community i, we get

$$f_i(\mathbf{x}) = (1 - x_i)p_i(\mathbf{x}) - x_i p_i(\mathbf{1} - \mathbf{x}) \tag{12}$$

V.3 Asymptotic results

The Markov process will get absorbed in finite time almost surely, by one of the two absorbing states where every agent shares the same opinion. Yet, a QSD exists and we now determine the set $\mathcal{A}$.

Regardless of the value of p, the only invariant sets are equilibria, i.e. points such that $f(\mathbf{x}) = 0$. For $p < 1/4$, the set of equilibria $\mathcal{E}$ has 9 elements: $(0, 0), (1, 1), (\frac{1}{2}, \frac{1}{2})$,

$$\mathbf{x}^1(p) := \left(\frac{1+\sqrt{1-4p}}{2}, \frac{1-\sqrt{1-4p}}{2}\right), \mathbf{x}^2(p) := \left(\frac{1-\sqrt{1-4p}}{2}, \frac{1+\sqrt{1-4p}}{2}\right),$$

as well as 4 other equilibria we do not specify, because they are always unstable. We observe that $(\frac{1}{2}, \frac{1}{2})$ is unstable, while $\mathbf{x}^1(p)$ and $\mathbf{x}^2(p)$ are stable if and only if $p < \bar{p} = \frac{1}{4}\left(2 - \sqrt{2}\right) \approx 0.146$. Also, $\mathbf{0}$ and $\mathbf{1}$ are in $\mathbf{S_0}$, so, if $p < \bar{p}$,

$$\mathcal{A} = \{\mathbf{x}^1(p), \mathbf{x}^2(p)\}$$

We apply Theorems 1 and 2. By noting that for both $\mathbf{x}^1(p)$ and $\mathbf{x}^2(p)$, both coordinates are symmetric with respect to $\frac{1}{2}$, and the difference between the two communities is large (i.e. $\sqrt{1-4p}$, which is between 0.65 and 1), we can state:

Theorem 3 *If $p < \bar{p}$,*

(i) Absorption time is exponential in N

(ii) Before absorption and regardless of initial conditions

- *on aggregate, opinions are in the neighborhood of 50% for both candidates*
- *society is polarized with two communities largely supporting opposite candidates*

The predictions of Theorem 3 are illustrated with the simulations presented in Figures 1, 2 and 4, where we observe exponential absorption time, asymmetric equilibria around 88% and 12% of A voters (since $p = 0.1$), and an aggregate of around 50% votes for both A and B.

The intuition behind these stable asymmetric equilibria is the following: at such states, the probability of an A voter (resp. B voter) from community 1 switching to B (resp. to A) is the same as the probability of a B voter (resp. A voter) from community 2 switching to A (resp. to B). Thus, switches occur in equal proportions in both communities but in opposite directions, keeping the aggregate proportions of A voters steady. To understand why communities with a high proportion of A voters do not drift to 100% A voters, note that there is a small probability (p) that these voters will meet with a high proportion of B voters. Thus there is always a chance that they switch to B.

Once the process gets very close to one of the two stable interior states, it takes a very large number of switches to reach consensus, since one of the two communities has to switch from a majority of A (resp. B) voters to a majority of B (resp. A) voters. These events happen with low probability, because the members of a given community are essentially interacting with like-minded voters. It thus requires a large number of events, each occurring with low probability, for the process to get absorbed into consensus.

When p is large, the set of stable equilibria is empty and the Markov process will get absorbed quickly. This is intuitive, since large p means high levels of interaction between communities. Therefore, as soon as one community starts having a high proportion of A voters, there is a high chance that a B voter from any community also switches to A. Once this process starts, it accelerates as the number of A voters increases, until the process is absorbed into consensus.

Finally, note that while $(\frac{1}{2}, \frac{1}{2})$ is also an equilibrium, it is nevertheless unstable for any value of p. Thus, even if initial conditions are set to 50% A voters in **both** communities, the process quickly drifts away. This is clearly illustrated in Figure 4, where initial conditions are close to fifty-fifty and opinions quickly drift away towards $\mathbf{x}^1(p)$ or $\mathbf{x}^2(p)$.

To further illustrate Theorem 2 and how deterministic approximation works, we present simulations in Figures 5 and 6. In Figure 5, we show trajectories of the process in a two-dimensional space where each dimension represents one community. These trajectories show two symmetric accumulation points, which correspond to $\mathbf{x}^1(p)$ and $\mathbf{x}^2(p)$. In Figure 6 we plot the vector field of the dynamical system $\dot{\mathbf{x}} = f(\mathbf{x})$, which approximates the trajectories of the random process when N gets large. As we observe, trajectories of the random process almost perfectly mirrors the flow of the dynamical system towards either $\mathbf{0}, \mathbf{1}$, $\mathbf{x}^1(p)$ and $\mathbf{x}^2(p)$, but ignores all trajectories leading to unstable equilibria.

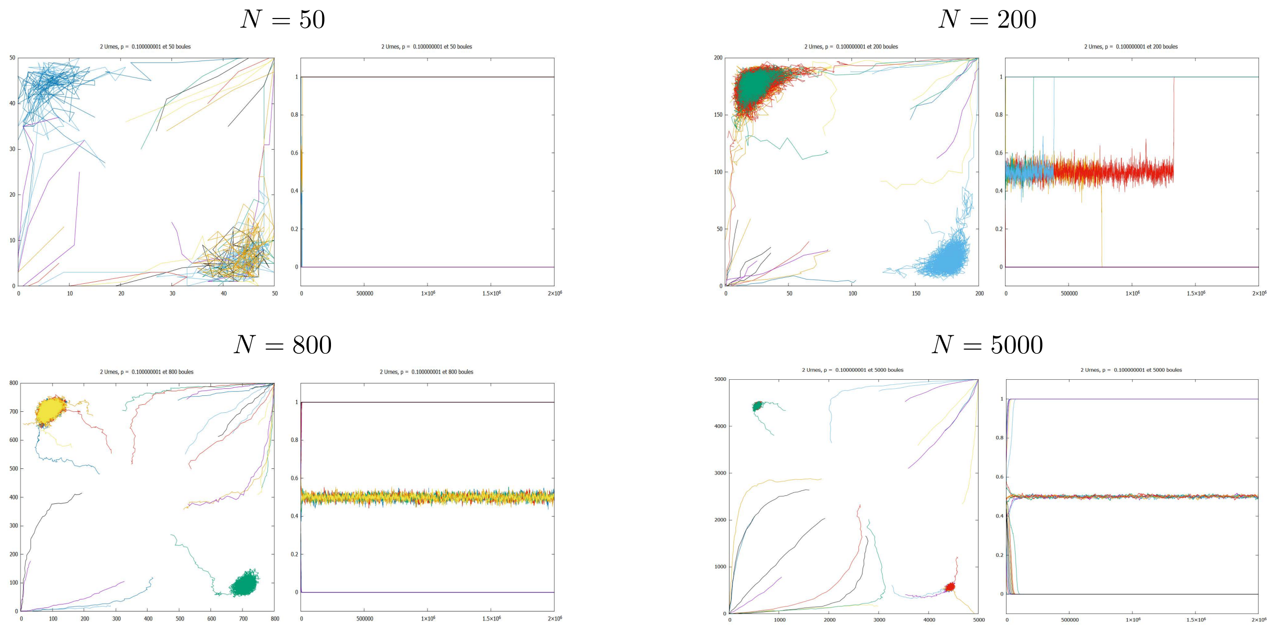

Figure 5: Each Panel depicts 24 realizations with random initial conditions and $p = 0.1$. **Left sub-panels:** the horizontal (resp. vertical) axis corresponds to the fraction of A-voters in community 1 (resp. community 2). The clusters in the upper left and lower right corners correspond to the support of the QSD. **Right sub-panel:** aggregate share of A-voters.

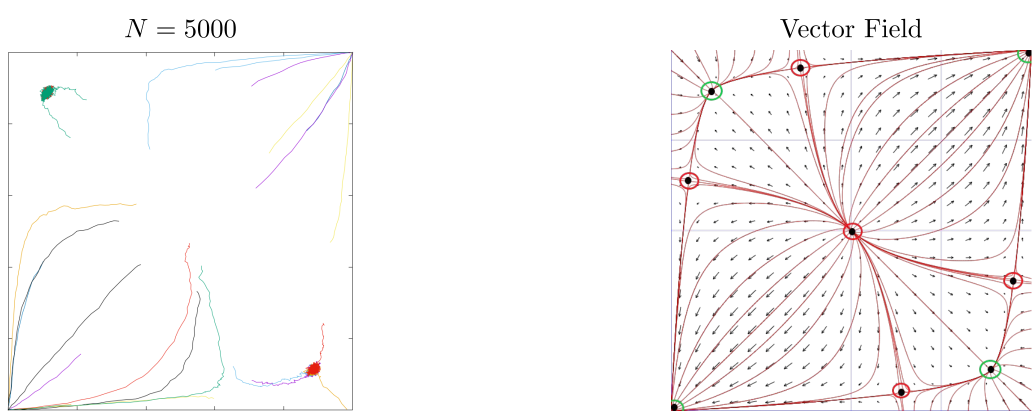

Figure 6: The left panel repeats the last simulation of Figure 5 with $N = 5000$ voters per community. The right panel shows the vector field of $\dot{\mathbf{x}}(t) = f(\mathbf{x}(t))$.

VI Application 2: Learning From Online Reviews

We now consider a model in the spirit of Acemoglu et al. (2022), where agents decide whether or not purchasing a given product, based on both their ex-ante valuation of the item and the reviews left by previous customers. We show that, despite good and bad products both disappearing, bad products disappear quickly while good products remain on the market for a very long time, and are correctly rated before disappearance.

VI.1 Description of the Model

We follow Acemoglu et al. (2022) for the main features of the model. A product has quality $Q \in \{0,1\}$. At each time $n \in \mathbb{N}^*$ a customer arrives and decides whether to purchase the product or not, after observing the current *rating* of the product, $R_n \in [0,1]$. He is endowed with a pair of private characteristics (θ_n, ζ_n), θ_n being his ex-ante valuation of the item, and ζ_n capturing his ex-post experience; formally, $(\theta_n)_n$ and $(\zeta_n)_n$ are sequences of i.i.d random variables with support respectively $[0,\overline{\theta}]$ and $\mathbb{R}$. Customer n takes his purchasing decision taking into account the rating he observes, R_n, and his individual characteristic θ_n: agent n buys the product if and only if:

$$D(R_n, \theta_n) \geq 0,$$

where $D : [0,1] \times [0,\overline{\theta}]$ has the following properties:

$(D1)$ $D(.,.)$ is strictly increasing in both arguments;

$(D2)$ there exists $\underline{R} \in]0,1[$ such that $D(\underline{R}, \overline{\theta}) = 0$.

The second assumption states that there is a rating $\underline{R}$ under which even the most enthusiastic customer decides not to buy the product, and the product disappears from the market. Given $(D1)$, this value is unique. Note also that $(D2)$ is the main difference with Acemoglu et al. (2022), where they assume that most optimistic customers will always buy the product.[15]

Customer n who buys the product then discovers its quality, and his ex-post experience of the product captured by ζ_n. A positive (resp. negative) ζ_n means the customer is positively (resp. negatively) surprised by its quality. He then leaves a review $r_n \in \{0,1\}$[16], depending on whether he was satisfied or not by the product. Again, this decision is threshold-based:

$$r_n = 1 \text{ iif } \rho(\theta_n, \zeta_n, Q) \geq 0,$$

where $\rho : [0,\overline{\theta}] \times \mathbb{R} \times \{0,1\}$ has the following properties:

$(R1)$ ρ is continuous and increasing in every argument;

$(R2)$ $\mathbb{P}\left[\rho(\overline{\theta}, \zeta, 1) \geq 0\right] > \underline{R} > \mathbb{P}\left[\rho(\overline{\theta}, \zeta, 0) \geq 0\right]$.

[15] In their paper, Acemoglu et al. (2022) focus on whether customers can learn the quality of the product. Thus they need that any product, even of poor quality, survives on the market and is purchased, regardless of how bad the reviews of previous customers are. This condition is necessary for learning, which could otherwise fail. In our adaptation, all products will disappear, but we focus on what happens before disappearance.

[16] This system can be made more complex: the reviews could take many values instead of just two, corresponding to different satisfaction levels.

To understand assumption ($R2$), remember that $\underline{R}$ is the *critical* rating under which no customer buys the product. On the other hand, $\mathbb{P}\left[\rho(\overline{\theta},\zeta,.)\geq 0\right]$ is the probability with which the most optimistic buyer leaves a positive review, depending on the quality of the product. It can also be interpreted as the expected value of the review that this customer will leave. If this value is greater (resp. smaller) than $\underline{R}$, it means that, on average, this customer thinks the product deserves a review greater (resp. lower) than the critical rating.

We now describe the process governing the rating of the product. The rating is an element of $[0,1]\cap\frac{1}{N}\mathbb{Z}$, where N represents the number of possible ratings. It increases (resp. decrease) by $\frac{1}{N}$ after receiving a series of positive (resp. decrease) reviews, with some probability $\nu(.)$.[17]

$$\text{If } r_n = 1,\ R^N_{n+1} = \begin{cases} R^N_n + \frac{1}{N} & \text{with probability } \nu(1-R^N_n), \\ R^N_n & \text{with probability } 1-\nu(1-R^N_n). \end{cases}$$

$$\text{If } r_n = 0,\ R^N_{n+1} = \begin{cases} R^N_n - \frac{1}{N} & \text{with probability } \nu(R^N_n), \\ R^N_n & \text{with probability } 1-\nu(R^N_n). \end{cases}$$

For simplicity in what follows we use $\nu(R) = R$.

VI.2 Existence of a QSD and Deterministic Approximation

Given $R > \underline{R}$, the event $\{D(R,\theta_n)\geq 0\}$ is non-null. Consequently, the map $R\mapsto d(R) := \mathbb{P}\left(D(R,\theta_n)\geq 0\right)$ is equal to zero on [0,$\underline{\text{R}}$], while it is strictly increasing on $]\underline{R},1]$. We can therefore define the map

$$R\in]\underline{R},1]\mapsto \pi(R) := \mathbb{P}\left[\rho(\theta_n,\zeta,Q)\geq 0 \mid D(R,\theta_n)\geq 0\right], \tag{13}$$

which, to a rating $R>\underline{R}$, associates the probability that a customer buying the product leaves a positive review. It is a strictly decreasing map on $]\underline{R},1]$[18], since lower θ types satisfy the $D(R,\theta)\geq 0$ condition as R increases, making the event $\{\rho(\theta_n,\zeta,Q)\geq 0\}$ less likely. Thus, for $R>\underline{R}$,

$$R^N_{n+1} - R^N_n = \frac{1}{N} Z(R^N_n)$$

where

$$Z(R) = \begin{cases} +1 & \text{with probability } (1-R)\pi(R)d(R), \\ -1 & \text{with probability } R(1-\pi(R))d(R), \\ 0 & \text{with probability } 1-(1-R)\pi(R)d(R) - R(1-\pi(R))d(R), \end{cases}$$

and when $R\leq\underline{R}$, $Z(R)=0$ with probability one since no customer buys the product[19]. We prove in the appendix (Lemma 2) that Assumption 2 holds in this model.

The drift associated to this random process is obtained by computing the expected value of $Z^N(R)$. In this case,

[17] Ratings do not necessarily change after each review. The probability $\nu(.)$ captures the idea that it takes more positive reviews to increase a rating when it is already high than when it is intermediate.

[18] This is the selection effect that Acemoglu et al. (2022) identify in their paper: as the rating gets higher, more customers buy the item, allowing for a higher probability that customers are disappointed and leave a bad review.

[19] To be rigourous we would need to make sure that R^N_n remains in $[0,1]$ by defining $Z^N(R)$ in such a way that $R=1$ whenever the increment $\frac{1}{N}$ causes R to exceed 1. We skip these details to keep the presentation simple.

$$f(R) = (1-R)\pi(R)d(R) - R(1-\pi(R))d(R) = (\pi(R)-R)d(R) \tag{14}$$

Since R will cross the boundary $\underline{R}$ at some point almost surely, the Markov process will get absorbed in finite time, and every product, whether of good or poor quality, will disappear from the market. Yet, a QSD exists and what happens to a product depends on the properties of the dynamical system $\dot{R} = f(R)$.

VI.3 Asymptotic results

We show that the process will be absorbed very fast if the product is of poor quality, while it will be absorbed very slowly when the product is of good quality, exhibiting radically different behaviors that cannot be detected by focusing only on steady-state.

Theorem 4 (Quasi-complete learning) *We have:*

- *If $Q = 1$, there is a unique value $R^* \in \mathbf{S}_* =]\underline{R}, 1]$ such that*

$$\pi(R^*) = R^*.$$

  The product will be on the market and rated approximately R^* for an exponentially long time (in N), before disappearance.

- *If $Q = 0$ the product quickly disappears from the market.*

We call this theorem "quasi-complete learning" to highlight the fact that, in the scenario of Acemoglu et al. (2022) where there is no complete learning, we can still predict that poor quality products will disappear very fast while good quality products will survive an extremely long time on the market with a high rating, before eventually disappearing.

Example 1 *Assume D and ρ are linear, as in Acemoglu et al. (2022): $D(R_n, \theta_n) = R_n + \theta_n - p$ and $\rho(\theta_n, \zeta_n, Q) = \theta_n + \zeta_n + Q - p$, where p is the price of the product. Assume also that $Q = 1$, θ is uniform on $[0, 1]$, ζ is uniform on $[-1/2, 1/2]$. Then, $\underline{R} = p - 1$ is the rating under which customers stop buying the product. Also, conditional on $\theta + R - p \geq 0$, θ is uniform on $[p - R, 1]$. Thus*

$$\pi(R) = \mathbb{P}\left(\theta + \zeta \geq p - 1 \mid \theta \geq p - R\right) = \mathbb{P}(X + \zeta \geq p - 1)$$

where $X \sim U[p - R, 1]$ and ζ are independent. We have[20]

$$\begin{aligned} \pi(R) &= \int_{p-R}^{1} \left(\int_{p-1-x}^{\frac{1}{2}} dz \right) \frac{1}{1-p+R} dx \\ &= \frac{1}{2}(4 - p - R) \end{aligned}$$

Then the QSD R^* satisfies $\pi(R^*; 1) = R^*$, i.e.

$$R^* = \frac{4-p}{3}$$

[20] The computations are true for every $p \geq \frac{3}{2}$, since otherwise the integral over z would need to be split into two parts. We omit these details for clarity of exposition.

We show some realizations of this random process for varying values of N in Figure 7, and show how the process settles around the QSD before being absorbed into disappearance. In this example, $p = 1.55$, so that the QSD is predicted to concentrate around $R^* = \frac{4-p}{3} \simeq 0.816$ and the threshold value is $\underline{R} = 0.55$. We intentionally start at $R_0 = 0.6$, which is quite close to the threshold value, to illustrate that concentration around the QSD is not due to favorable initial conditions.

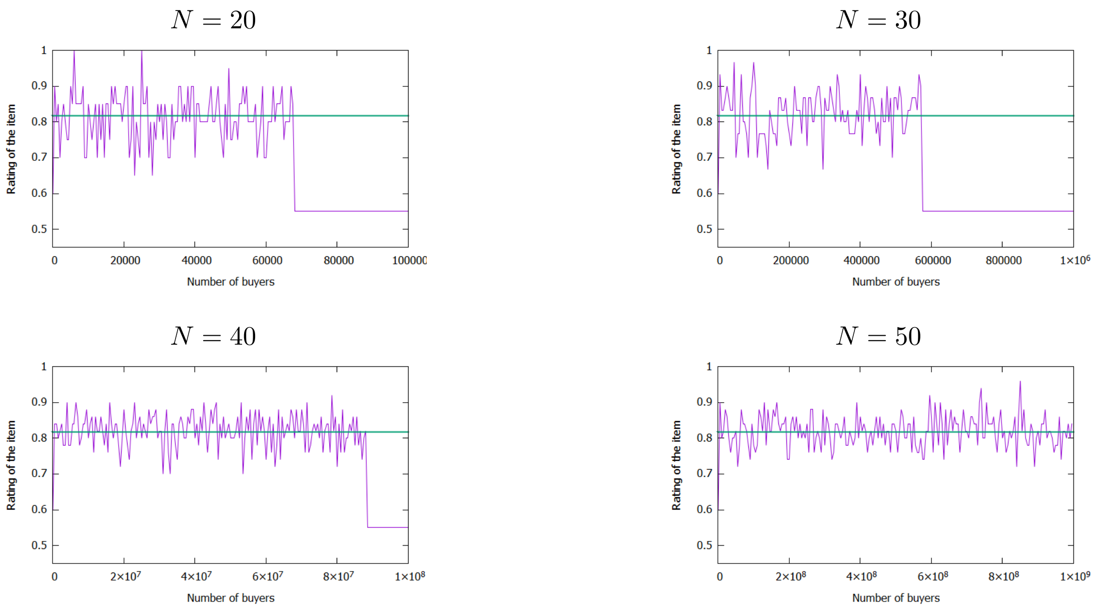

Figure 7: Rating of the product, R_n for N varying from 20 to 50, when $p = 1.55$, $\underline{R} = 0.55$, and $R_0 = 0.6$. The scale changes for each panel. In green is the line $R^* = 0.816$ around which the process oscillates before being absorbed.

Theorem 4 shows how platforms, who wish to avoid good product from disappearing ”by mistake” while making sure that bad products will disappear fast, should have a fine enough rating system. If reviews cannot cause the rating to jump too much, good products will be ”almost indefinitely” present on the platform.

VII Application 3: The Persistence of Superstitions and False Beliefs

Individuals form beliefs about a true state of nature by interpreting a finite number of “facts.” Each fact can be interpreted as supporting either the true state (e.g. “Earth is round”) or the false one (e.g. “Earth is flat”), and these interpretations are revised over time. We show that false beliefs - although ultimately transient - can persist for very long periods of time before disappearing.

VII.1 Description of the model

A belief is constructed from N facts. Each fact can be understood in two possible ways: either as supporting the true state (assigned value +1) or the false one (assigned value 0). For example, the statement "We cannot reach the edge of the Earth" can be interpreted as "because the Earth is surrounded by an ice wall" (supporting flat Earth, value 0), or as "because a sphere has no edge" (supporting round Earth, value +1). At period n, the opinion of the agent is therefore represented by

$$q_n^N \in \left\{ \tfrac{k}{N} : k = 0, \ldots, N \right\},$$

the fraction of facts currently interpreted as +1. The parameter N can thus be interpreted as the dimension of the belief space, and hence as a measure of the complexity of the problem faced by the agent.

The belief evolves as follows. At each period, the agent randomly selects one fact and re-examines it. Its interpretation may change depending on the overall state of belief: the agent has a tendency to align facts with the dominant interpretation, reflecting a preference for internal consistency. In addition, Nature occasionally intervenes, forcing the selected fact to take the value +1, i.e., providing an unquestionable signal in favor of the true state. Crucially, Nature never induces a switch from +1 to 0, which introduces a strong asymmetry in the dynamics.

The two extreme beliefs $q = 0$ and $q = 1$ correspond to complete certainty in the false and true state, respectively. Because Nature can always transform some 0's into +1's, the state $q = 0$ is not absorbing. In contrast, once $q = 1$ is reached, all facts are interpreted as +1 and the belief remains there forever: $q = 1$ is an absorbing state

Formally, the dynamics can be written

$$q_{n+1}^N = q_n^N + \frac{1}{N} Z(q_n^N),$$

where the random variable $Z(q)$[21] describes the change induced by reinterpreting the questioned fact:

$$Z(q) = \begin{cases} +1 & \text{w.p.} \quad (1-q)\left(q + \frac{1}{N}\right)^2, \\ -1 & \text{w.p.} \quad q(1-q)^2, \\ 0 & \text{w.p.} \quad 1 - \mathbb{P}(Z(q) \in \{+1, -1\}). \end{cases} \tag{15}$$

The intuition is straightforward. For $Z(q) = +1$, the selected fact must currently be 0 (probability $1 - q$), and then be reinterpreted as +1. This reinterpretation occurs if the fact is compared with two reference facts carrying the value +1, which happens with probability $(q + 1/N)^2$ (where $1/N$ accounts for Nature's intervention). Conversely, for $Z(q) = -1$, the questioned fact must be +1 (probability q) and both reference facts must be 0 (probability $(1-q)^2$). Note that, in this case, Nature does not intervene: a +1 cannot be turned into a 0 by an external signal. Finally, $Z(q) = 0$ corresponds to situations where the fact keeps its current interpretation.

VII.2 Existence of a QSD and Deterministic Approximation

The process is a Markov chain and computing the expected value of Z we get

$$f(q) = (1-q)q^2 - q(1-q)^2,$$

[21] In the model we present in Section III, the switching kernels only depend on $\mathbf{x}$. Everything works when they also depend on N, as long as there exist a limit kernel μ such that $|\mu^N(\mathbf{d} \mid \mathbf{x}) - \mu(\mathbf{d} \mid \mathbf{x})| \leq \frac{K}{N}$, where $K > 0$ does not depend on $\mathbf{d}$ and $\mathbf{x}$.

As a consequence, the behavior of the QSD will depend on the behavior of the system

$$\dot{q} = f(q)$$

VII.3 Asymptotic results

The deterministic system has three stationary points, $q = 0$, $q = 1$ and $q = \frac{1}{2}$. Only 0 and 1 are stable stationary points, and only 0 is in $\mathcal{A}$. Thus:

Theorem 5 *We have:*

- *Absorption time into* $q = 1$ *is exponential in* N*, the complexity of the problem*
- *The quasi-stationary distribution is concentrated around* $q = 0$

In this application again, we have used a simple model for which the long run conclusion does not fit what we try to explain (the persistence of false beliefs). Instead of twisting this model to make individual behavior more complex and allow for the emergence of persistence, we look at pre-convergence behavior, and find that this might explain why we observe some individuals still firmly believing earth is flat.

As for the previous applications, we present some simulations of the process.

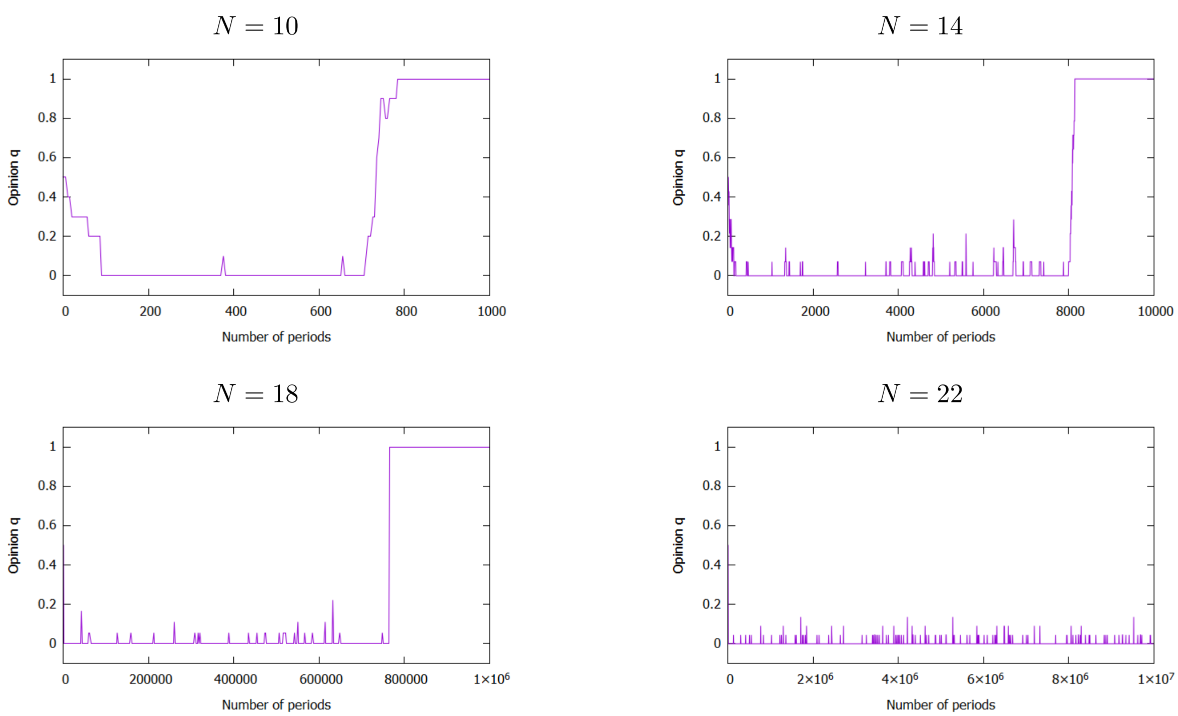

Figure 8: Beliefs for N varying from 10 to 22. The scale changes for each panel, as absorption time gets larger.

VIII An Example Where the QSD is Not Relevant

Let us consider application 3, with another updating process where there is no inertia , i.e. the first fact to which the agent compares the questioned fact is decisive. In that case, the process writes

$$q_{n+1}^N = q_n^N + \frac{1}{N} Z(q_n^N)$$

where

$$Z(q) = \begin{cases} 1 & \text{w.p.} \quad (1-q)\left(q + \frac{1}{N}\right), \\ -1 & \text{w.p.} \quad q(1-q), \\ 0 & \text{w.p.} \quad 1 - \mathbb{P}\left(Z(q) \in \{1, -1\}\right) \end{cases} \tag{16}$$

As previously, this process has only one absorbing state which is $q = 1$, (so that assumption 2 holds with $\mathbf{S}_0 = \{1\}$), and this process satisfies all the requirements for a QSD to exist. However this QSD is irrelevant to understand the behavior of the process.

Indeed, computing the expected value of $Z(q)$, we find $f \equiv 0$. Therefore, the drift is null and $\mathcal{A} = \emptyset$, so that Theorem 1 and 2 do not apply. As we show now, absorption time will be short, and behavior before absorption will be random, uniformly distributed on $[0, 1]$.

In this particular setting, we can compute the QSD by hand. The transition matrix restricted to transient states, $\mathbf{Q}_N$, is a tridiagonal matrix of dimension $N \times N$, where states vary from 0 to $\frac{N-1}{N}$. By observing that every column of $\mathbf{Q}_N$ sums up to $1 - \frac{1}{N^2}$, we deduce that $\pi_N \mathbf{Q}_N = \lambda_N \pi_N$ holds with

$$\pi_N = \left(\frac{1}{N}, ..., \frac{1}{N}\right); \;\; \lambda_N = 1 - \frac{1}{N^2}.$$

As we see, the time spent in each state by the process before absorption is uniformly distributed on every state from 0 to $\frac{N-1}{N}$. Further, expected absorption time can be computed to be of the order of N^2. In consequence, the process will be quickly absorbed and there is no remarkable information to extract on the pre-absorption behavior.

These two points are illustrated in the following simulations.

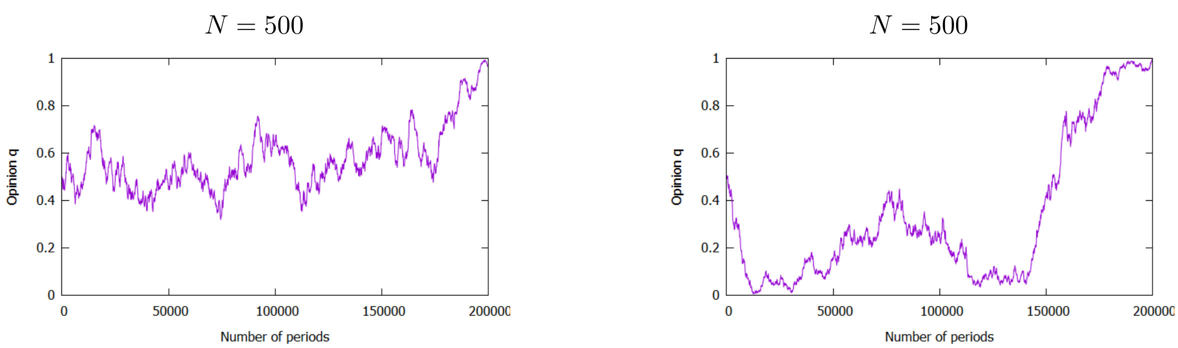

Figure 9: Belief for $N = 500$. It behaves erratically, never stabilizes, and gets absorbed rapidly.

References

Acemoglu, D., Makhdoumi, A., Malekian, A., and Ozdaglar, A. (2022). Learning from reviews: The selection effect and the speed of learning. *Econometrica*, 90(6):2857–2899.

Acemoglu, D. and Ozdaglar, A. (2011). Opinion dynamics and learning in social networks. *Dynamic Games and Applications*, 1(1):3–49.

Bartlett, M. S. (1957). On theoretical models for competitive and predatory biological systems. *Biometrika*, 44(1/2):27–42.

Bartlett, M. S. (1960). Stochastic population models; in ecology and epidemiology. Technical report.

Benaïm, M. and Weibull, J. W. (2003). Deterministic approximation of stochastic evolution in games. *Econometrica*, 71(3):873–903.

Binmore, K. and Samuelson, L. (1994). An economist's perspective on the evolution of norms. *Journal of Institutional and Theoretical Economics (JITE)/Zeitschrift für die gesamte Staatswissenschaft*, pages 45–63.

Binmore, K. G., Samuelson, L., and Vaughan, R. (1995). Musical chairs: Modeling noisy evolution. *Games and economic behavior*, 11(1):1–35.

Compte, O. (2021). Belief formation and the persistence of superstitions. *Working paper.*

Conley, C. C. (1978). *Isolated invariant sets and the Morse index.* Number 38. American Mathematical Soc.

Darroch, J. N. and Seneta, E. (1965). On quasi-stationary distributions in absorbing discrete-time finite markov chains. *Journal of Applied Probability*, 2(1):88–100.

Dupuis, P. and Ellis, R. S. (2011). *A weak convergence approach to the theory of large deviations*, volume 902. John Wiley & Sons.

Ellison, G. (1993). Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–1071.

Faure, M., Schreiber, S. J., et al. (2014). Quasi-stationary distributions for randomly perturbed dynamical systems. *The Annals of Applied Probability*, 24(2):553–598.

Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica: Journal of the Econometric Society*, pages 29–56.

Keynes, J. M. (1923). *A tract on monetary reform.* London, Macmillan.

Kreindler, G. E. and Young, H. P. (2013). Fast convergence in evolutionary equilibrium selection. *Games and Economic Behavior*, 80(C):39–67.

Méléard, S., Villemonais, D., et al. (2012). Quasi-stationary distributions and population processes. *Probability Surveys*, 9:340–410.

Ruelle, D. (1981). Small random perturbations of dynamical systems and the definition of attractors. *Communications in Mathematical Physics*, 82(1):137–151.

Yaglom, A. M. (1947). Certain limit theorems of the theory of branching random processes. In *Doklady Akad. Nauk SSSR (NS)*, volume 56, page 3.

Young, H. P. (1993). The evolution of conventions. *Econometrica: Journal of the Econometric Society*, pages 57–84.

Zeitouni, A. D. O. and Dembo, A. (1998). Large deviations techniques and applications. *Applications of Mathematics*, 38.

IX Appendix A: Quasi-stationary distributions

IX.1 Existence, uniqueness, absorption time

Here we compile standard general results from Darroch and Seneta (1965) on QSDs for finite state Markov chains. A thorough exposition can be found in Méléard et al. (2012).

An alternative formulation of this definition is the following. The transition matrix of the Markov process, $(\mathbf{P}(\mathbf{x}, \mathbf{x}'))_{\mathbf{x},\mathbf{x}' \in \mathbf{E}}$, can be written as

$$\mathbf{P} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{R} & \mathbf{Q} \end{bmatrix},$$

where $\mathbf{I}$ is the $|\mathbf{E}_0| \times |\mathbf{E}_0|$ identity matrix, $\mathbf{0}$ is an $|\mathbf{E}_0| \times |\mathbf{E}_*|$ matrix of zeroes, $\mathbf{R}$ is an $|\mathbf{E}_*| \times |\mathbf{E}_0|$ matrix whose entries are the transition probabilities from transient states to absorbing states. In what follows we restrict our attention to cases where $\mathbf{R}$ is not the null matrix, meaning that there exists at least one state in the transient space from which the system can be absorbed in one step. Finally $\mathbf{Q}$ is the restriction of $\mathbf{P}$ to the transient space. A QSD can alternatively be defined as a probability measure π on $\mathbf{E}_*$ such that

$$\pi \mathbf{Q} = \lambda \pi \tag{17}$$

with $\lambda := \mathbb{P}_\pi(T > 1)$, the probability that the process is not absorbed in one step conditional on being initially distributed as π. In other words, if π is a QSD for $(\mathbf{x}_n)_{n \geq 0}$ then it is a left eigenvector associated with the matrix $\mathbf{Q}$.

Theorem A (Darroch and Seneta (1965)). *Suppose that* $\mathbf{Q}$ *is irreducible and aperiodic. Then* $(\mathbf{x}_n)_n$ *admits a unique quasi-stationary distribution* π*. Moreover,*

1) π *has full support on* $\mathbf{E}_*$*:* $\pi(\mathbf{x}) > 0$ *for all* $\mathbf{x} \in \mathbf{E}_*$*;*

2) *We have* $\lambda = \mathbb{P}_\pi(T > 1) \in (0, 1)$*;*

3) *If initial conditions are distributed as* π*, then* T *is geometrically distributed:*

$$\mathbb{P}_\pi(T > k) = e^{-k\, log(\frac{1}{\lambda})}, \quad and \quad \mathbb{E}_\pi[T] = \frac{1}{log(\frac{1}{\lambda})}.$$

4) *We have the ergodic property:*

$$\lim_{n\to+\infty} \mathbb{P}_{\mathbf{x}}\left(\mathbf{x}_n \in \Gamma \mid T > n\right) = \pi(\Gamma), \;\; \forall \mathbf{x} \in \mathbf{E}_*, \; \forall \Gamma \subseteq \mathbf{E}_*. \tag{18}$$

Since absorption occurs with positive probability from some transient state, the submatrix $\mathbf{Q}$ is (strictly) substochastic, implying that its spectral radius $\rho(\mathbf{Q})$ is strictly smaller than one. The first two points are thus direct consequences of the Perron-Frobenius theorem: $\rho(\mathbf{Q})$ is a simple eigenvalue of $\mathbf{Q}$, associated with a positive eigenvector. Moreover any non-negative eigenvector of $\mathbf{Q}$ is necessarily associated with $\rho(\mathbf{Q})$. By (17), π is a non-negative eigenvector, hence it is associated with $\rho(\mathbf{Q})$, and necessarily $\lambda = \rho(\mathbf{Q}) \in (0,1)$.

In the context of interest here, this theorem conveys two principal messages, captured by points 3 and 4. First, the absorption time depends on how large λ is. Recall that λ can be interpreted as the probability that the process does not get absorbed after one step. If λ is "much smaller"[22] than one, the process will be absorbed in a few steps, and therefore analyzing its QSD is irrelevant. However, if λ is "very close" to one, the absorption time will be "extremely long", and analyzing its QSD becomes critical to understand the economic phenomena under scrutiny.

Second, note that according to (3), a process distributed as π will remain distributed as π, as long as it remains in the transient space. However, if the process was never distributed as π in the first place, the non-absorbed process could look totally different from the QSD. This theorem rules out that possibility since point 4 says that pre-absorption behavior does not depend on initial conditions. It guarantees that, if absorption time is "long", the process will indeed behave like its QSD, making it the relevant outcome to analyze.

Finally, it is important to note that there is no general mathematical tool to compute the QSD. In small societies, finding the eigenvalue λ, the corresponding eigenvector π and therefore explicitely computing the QSD is feasible. However, in economics we are generally interested in situations where systems are "large", in which case explicitely computing the QSD is impossible.

X Appendix B: Proofs

X.1 Proof of Theorem 1.

We prove the results under the more general assumption that ϕ admits an attractor[23] A in the interior of $\mathbf{S}_*$. Given a set B, let $N^\delta(B) := \{\mathbf{x} \in \mathbf{S} : \; d(\mathbf{x}, B) < \delta\}$. Let A be an attractor included in $\mathbf{S}_*$. Then there exists an open neighborhood U of A such that $\phi(\mathrm{cl}(U), 1) \subseteq U$, where $\mathrm{cl}(U)$ is the closure of set U (see Conley (1978) or Ruelle (1981) for instance). Since $\phi(\mathrm{cl}(U), 1)$ is compact and U is open, there exists $\delta > 0$ such that $N^\delta(\phi(\mathrm{cl}(U), 1)) \subseteq U$. By definition of π_N, we have

$$\begin{aligned}
\lambda_N^N \pi_N\left(\mathrm{cl}(U)\right) &= \sum_{\mathbf{x}\in\mathbf{S}_*^N} \pi_N(\mathbf{x}) \mathbb{P}\left(\mathbf{x}_N^N \in \mathrm{cl}(U) \mid \mathbf{x}_0^N = \mathbf{x}\right) \\
&\geq \sum_{\mathbf{x}\in\mathrm{cl}(U)} \pi_N(\mathbf{x}) \mathbb{P}\left(\mathbf{x}_N^N \in \mathrm{cl}(U) \mid \mathbf{x}_0^N = \mathbf{x}\right) \\
&\geq \pi_N(\mathrm{cl}(U)) \inf_{\mathbf{x}\in\mathrm{cl}(U)} \mathbb{P}\left(\mathbf{x}_N^N \in \mathrm{cl}(U) \mid \mathbf{x}_0^N = \mathbf{x}\right)
\end{aligned}$$

[22] All terms between quotes in this paragraph will be made precise in the next sections.

[23] The set A is an *attractor* for the semi-flow ϕ if there exists an open neighborhood U of A such that $\lim_{t\to+\infty} \sup_{\mathbf{x}\in U} d\left(\phi(\mathbf{x}, t), A\right) = 0$.

Since $\phi(\mathbf{x},1) \in \phi(\mathrm{cl}(U),1)$ for any $\mathbf{x} \in \mathrm{cl}(U)$, the event $\left\{\mathbf{x}_N^N \notin \mathrm{cl}(U)\right\}$ is contained in the event $\left\{\|\phi(\mathbf{x},1) - \mathbf{x}_N^N\| \geq \delta\right\}$, conditional to $\mathbf{x}_0^N = \mathbf{x}$. Hence

$$\begin{aligned}\sup_{\mathbf{x}\in\mathrm{cl}(U)} \mathbb{P}\left[\mathbf{x}_N^N \notin \mathrm{cl}(U) \mid \mathbf{x}_0^N = \mathbf{x}\right] &\leq \sup_{\mathbf{x}\in\mathrm{cl}(U)} \mathbb{P}\left[\|\phi(\mathbf{x},1) - \mathbf{x}_N^N\| \geq \delta \mid \mathbf{x}_0^N = \mathbf{x}\right] \\ &\leq 2ke^{-c\delta^2 N},\end{aligned}$$

where the exponential bound can be proved exactly as Lemma 1 in Benaïm and Weibull (2003)[24]. Thus we have

$$\lambda_N^N \geq 1 - \sup_{\mathbf{x}\in\mathrm{cl}(U)} \mathbb{P}\left(\mathbf{x}_N^N \notin \mathrm{cl}(U) \mid \mathbf{x}_0^N = \mathbf{x}\right) \geq 1 - 2ke^{-c\delta^2 N}, \tag{19}$$

which implies that

$$\lambda_N \geq \exp\left(\frac{1}{N}\ln\left(1 - 2ke^{-c\delta^2 N}\right)\right) \geq \exp\left(-\frac{1}{N}2ke^{-c\delta^2 N}\right) \sim_{N\to+\infty} 1 - \frac{1}{N}2ke^{-c\delta^2 N},$$

providing the desire result, by choosing $0 < C_\lambda < c\delta^2$. ■

X.2 Proof of Theorem 2

We want to prove that the support of the limiting measure is contained in the attractors of the mean dynamics. The key tool in order to relate the QSD of our process π_N to the semi-flow $(\phi(\mathbf{x},t))_{\mathbf{x}\in[0,1]^k, t\geq 0}$ is proving that the trajectory of the process remains close to the solution curves of the flow with high probability. However we want more than this, in the sense that we want to prove that the relevant states are the stable equilibria of the mean dynamics, and only those. Hence we also need to prove that it is reasonably likely that the random process does not spend too much time around an unstable equilibrium; more generally we need to show that the random process has a reasonable chance to wander away from the deterministic flow.
We adapt a result by Faure et al. (2014) to our settings. We state this result in Section X.2.1 (Proposition 1), we show that our model fits in this framework in Section X.2.2, and check that the necessary conditions are fulfilled in Section X.2.3.

X.2.1 Limiting measures of QSDs for random perturbations

Given a continuous map $F : \mathbf{S} \to \mathbf{S}$ such that $F(\mathbf{S}_0) \subseteq \mathbf{S}_0$ and $F(\mathbf{S}_*) \subseteq \mathbf{S}_*$, a *random perturbation* of F is a family of Markov chains $\left\{\mathbf{X}^N\right\}$ on $\mathbf{S}$ such that

- $\mathbf{S}_0$ is absorbing for $\mathbf{X}^N$: $\mathbb{P}\left(\mathbf{X}_1^N \in \mathbf{S}_* \mid \mathbf{X}_0^N = \mathbf{x}\right) = 0$, for all $\mathbf{x} \in \mathbf{S}_0$;
- for any $\delta > 0$,
$$\lim_{N\to+\infty} \sup_{\mathbf{x}\in\mathbf{S}} \mathbb{P}\left(\mathbf{X}_1^N \notin N^\delta(F(\mathbf{x})) \mid \mathbf{X}_0^N = \mathbf{x}\right) = 0.$$

Suppose that all the invariant sets of the dynamics (except for the absorbing set) are contained in an open set whose closure is contained in $\mathbf{S}_*$. Let K be such a compact set

[24] The only difference between their settings and ours is the state space, as the state space is the simplex in Benaïm and Weibull (2003). Hence since we need this inequality to hold in a compact set in the interior of $\mathbf{S}$, exactly the same proof using exponential martingale inequalities will lead to the same conclusion.

Assumption 3 (Large deviation assumption) *There exists a map*

$$\rho : (\mathbf{x}, \mathbf{y}) \in \mathbf{S}_* \times \mathbf{S}_* \to [0, +\infty]$$

such that the following holds:

- ***Regularity of the rate function:*** $\forall\, \mathbf{x} \in \mathbf{S}_*$ *and* $T > 0$,
  - $\rho(\mathbf{x}, \mathbf{y}) = 0$ *if and only if* $\mathbf{y} = F(\mathbf{x})$
  - *the map* $\mathbf{y} \mapsto \rho(\mathbf{x}, \mathbf{y})$ *has compact level sets:*

$$\{\mathbf{y} \in \mathbf{S}_* : \ \rho(\mathbf{x}, \mathbf{y}) \leq L\} \text{ is compact, for any } L \geq 0;$$

- ***Large deviation lower bound:*** *for any open set* $\mathbf{U} \subseteq \mathbf{S}_*$,

$$\liminf_{N \to +\infty} \frac{1}{N} \log \mathbb{P}\left(\mathbf{X}_1^N \in \mathbf{U} \mid \mathbf{X}_0^N = \mathbf{x}\right) \geq - \inf_{\mathbf{y} \in \mathbf{U}} \rho(\mathbf{x}, \mathbf{y}) \tag{20}$$

  uniformly in $\mathbf{x} \in K$;

- ***Large deviation upper bound:*** *for any closed set set* $\mathbf{C} \subseteq \mathbf{S}_*$,

$$\limsup_{N \to +\infty} \frac{1}{N} \log \left(\mathbf{X}_1^N \in \mathbf{C} \mid \mathbf{X}_0^N = \mathbf{x}\right) \leq - \inf_{\mathbf{y} \in \mathbf{C}} \rho(\mathbf{x}, \mathbf{y}), \tag{21}$$

  uniformly in $\mathbf{x} \in K$.

Assumption 4 (Likely absorption near $\mathbf{S}_0$) *Given* $c > 0$, *there exists an open neighborhood* U_c *of* $\mathbf{S}_0$ *such that*

$$\liminf_{N \to +\infty} \inf_{\mathbf{x} \in U_c} \frac{1}{N} \log \mathbb{P}_{\mathbf{x}}\left(\mathbf{X}_1^N \in \mathbf{S}_0 \mid \mathbf{X}_0^N = \mathbf{x}\right) \geq -c. \tag{22}$$

Remark 1 *Note that inequality (22) cannot be obtained as a corollary of (20) because the interior of* $\mathbf{S}_0$ *might be empty.*

Proposition 1 (Theorem 2.7 in Faure et al. (2014)) *Assume that Assumptions 3 and 4 hold. Then any limiting measure of* $(\pi_N)_N$ *is supported by the linearly stable equilibria of the flow.*

X.2.2 A random perturbation associated to $\mathbf{x}^N$.

Define the auxiliary discrete-time process $\mathbf{X}_t^N := \mathbf{x}_{tN}^N$, for $t \in \mathbb{N}^*$. Also introduce $F : \mathbf{S} \to \mathbf{S}$ as $F(\mathbf{x}) := \phi(\mathbf{x}, 1)$. Then

$$\mathbb{P}\left(\mathbf{X}_1^N \in \Gamma \mid \mathbf{X}_0^N = \mathbf{x}\right) = \mathbb{P}\left(\mathbf{x}_N^N \in \Gamma \mid \mathbf{x}_0^N = \mathbf{x}\right), \ \forall \text{ Borel sets in } \mathbf{S}.$$

As we already stated in the proof of Theorem 1,

$$\mathbb{P}\left(\mathbf{X}_1^N \notin N^\delta(F(\mathbf{x})) \mid \mathbf{X}_0^N = \mathbf{x}\right) = \mathbb{P}\left[\|F(\mathbf{x}) - \mathbf{X}_1^N\| \geq \delta \mid \mathbf{X}_0^N = \mathbf{x}\right] \leq 2ke^{-c\delta^2 N},$$

implying that the family $\{\mathbf{X}^N\}_N$ is a random perturbation of F. Also, for $\mathbf{x} \in \mathbf{S}_0$, we have $f(\mathbf{x}) = \mathbf{0}$. Consequently $F(\mathbf{x}) = \mathbf{x}$; also, for any $\mathbf{x} \in \mathbf{S}_*$, we necessarily have $\phi(\mathbf{x}, t) \in \mathbf{S}_*$ for any $t \geq 0$, by continuity of μ.

Thus $F(\mathbf{x}) \in \mathbf{S}_*$. The Markov chain $\mathbf{X}^N$ inherits the absorbing property of $\mathbf{x}^N$: $\mathbb{P}(\mathbf{X}_1^N \in \mathbf{S}_* \mid \mathbf{X}_0^N = \mathbf{x}) = 0$, for all $\mathbf{x} \in \mathbf{S}_0$.

Remembering that $\mathbf{Q}_N$ is the transition of $\mathbf{x}^N$ restricted to $\mathbf{S}_*$, we have

$$\mathbb{P}_{\pi_N}\left(\mathbf{X}_1^N \in \Gamma\right) = \mathbb{P}_{\pi_N}\left(\mathbf{x}_N^N \in \Gamma\right) \pi_N \mathbf{q}_N^N(\Gamma) = \lambda_N^N \pi_N.$$

Thus the QSD of the Markov chain $\mathbf{X}^N$ is also π_N, and the associated eigenvalue is λ_N^N. Assumption 4 holds for $(\mathbf{X}_1^N)_N$, by Assumption 2. Proving Theorem 2 thus amounts to checking that Assumption 3 holds for the family $(\mathbf{X}_1^N)_N$. We proceed to showing this in the next subsection.

X.2.3 Derivation of Hypothesis 3

We use Theorem 6.3.3 in Dupuis and Ellis (2011) to derive the large deviation principle of Assumption 3. Consider a family of $\mathbb{R}^k$-valued random sequences $\{\chi^N\}_{N\geq 1}$ satisfying the recursive formula

$$\chi_{n+1}^N = \chi_n^N + \frac{1}{N} Z(\chi_n^N),$$

where, conditional to $\chi_n^N = \mathbf{x}$, Z is distributed accordingly to $\nu(\cdot \mid \mathbf{x})$. Also call $(\chi^N(t))_{t\in[0,1]} \in \mathcal{C}\left([0,1], \mathbb{R}^k\right)$ the piecewise linear interpolated process associated to $(\chi_n^N)_n$:

$$\chi^N(t) := \chi_j^N + (Nt - j)\,(\chi_{j+1}^N - \chi_j^N)\ \ \forall t \in [j/N, (j+1)/N]\,,\ \ j = 0, ..., N-1.$$

Now given $\mathbf{x} \in \mathbb{R}^k$, let $\mathbf{I}(\mathbf{x}, \cdot)$ be defined on $\mathcal{C}([0,1], \mathbb{R}^k)$ as

$$\mathbf{I}(\mathbf{x}, \varphi) := \begin{cases} \int_0^1 L(\varphi(t), \dot{\varphi}(t))dt & \text{if } \varphi \text{ is absolutely continuous and } \varphi(0) = \mathbf{x}, \\ +\infty & \text{otherwise} \end{cases} \tag{23}$$

where L is defined as follows:

$$L\left(\mathbf{x}, \mathbf{v}\right) := \sup_{\mathbf{z}\in\mathbb{R}^k} \left\{ \langle \mathbf{z}, \mathbf{v} \rangle - \log \int e^{<\mathbf{z},\mathbf{y}>} \nu(d\mathbf{y} \mid \mathbf{x}) \right\}.$$

The following theorem provides sufficient conditions on the stochastic kernel $\nu\left(\cdot \mid \mathbf{x}\right)\right)_{\mathbf{x}\in\mathbb{R}^k}$ under which the sequence of interpolated processes satisfies a uniform large deviation principle.

Theorem 6.3.3 in Dupuis and Ellis (2011) *Suppose that the family of measures $\{\nu(\cdot \mid \mathbf{x})\}_{\mathbf{x}\in\mathbb{R}^k}$ satisfies the following set of assumptions*

(i) $\sup_{\mathbf{x}\in\mathbb{R}^k} \int_{\mathbb{R}^k} \exp <\mathbf{z}, \mathbf{y}> \nu(d\mathbf{y} \mid \mathbf{x}) < +\infty$ *for any* $\mathbf{z} \in \mathbb{R}^k$*;*

(ii) the map $\mathbf{x} \in \mathbb{R}^k \mapsto \nu(\cdot \mid \mathbf{x}) \in \mathcal{P}(\mathbb{R}^k)$ *is continuous for the weak convergence topology on* $\mathbb{R}^k$*;*

(iii) the set $ri\,(co(Supp(\nu(\cdot, \mathbf{x}))))$[25] *is independent of* $\mathbf{x}$ *and contains* $\mathbf{0}$.

[25] The relative interior of the convex hull of μ's support.

Then we have

- *for any* $\mathbf{x} \in \mathbb{R}^k$,
    - $\mathbf{I}(\mathbf{x}, \varphi) = 0$ *if and only if* $\varphi(t) = \psi(\mathbf{x}, t)\ \forall t \in [0, T]$, *where* ψ *is the flow associated with the kernel* ν*;*
    - *the map* $\varphi \in \mathcal{C}\left([0,1], \mathbb{R}^k\right) \mapsto \mathbf{I}(\mathbf{x}, \varphi)$ *has compact level sets.*
- *for any open set* $\mathcal{U} \subseteq \mathcal{C}([0,1], \mathbb{R}^k)$,

$$\liminf_{N \to +\infty} \frac{1}{N} \log \mathbb{P}\left((\chi^N(t))_{t\in[0,1]}) \in \mathcal{U} \mid \chi^N(0) = \mathbf{x}\right) \geq - \inf_{\varphi \in \mathcal{U}} \mathbf{I}(\mathbf{x}, \varphi),$$

- *for any closed set* $\mathcal{C} \subseteq \mathcal{C}([0,1], \mathbb{R}^k)$

$$\limsup_{N \to +\infty} \frac{1}{N} \log \mathbb{P}\left((\chi^N(t))_{t\in[0,1]} \in \mathcal{C} \mid \chi^N(0) = \mathbf{x}\right) \leq - \inf_{\varphi \in \mathcal{C}} \mathbf{I}\left(\mathbf{x}, \varphi\right)$$

uniformly for $\mathbf{x} \in K$.

Define $\nu(\cdot, \mathbf{x})$ as follows

$$\nu(\cdot \mid \mathbf{x}) := \begin{cases} \mu(\cdot \mid \mathbf{x}) & \text{if } \mathbf{x} \in \mathbf{S}_* \\ \mu(\cdot \mid \Pi_{\mathbf{S}_*}(\mathbf{x})) & \text{otherwise} \end{cases}$$

where $\Pi_{\mathbf{S}}(\mathbf{x}))$ is the orthogonal projection of $\mathbf{x}$ on $\mathbf{S}_*$. We now check that $\nu(\cdot \mid \mathbf{x})$ satisfies the assumptions of previous theorem: point (i) holds because, for any $\mathbf{x} \in \mathbb{R}^k$, $\nu(\cdot \mid \mathbf{x})$ has finite support. Property (ii) directly follows from the continuity of $\mathbf{x} \mapsto \mu(\cdot \mid \mathbf{x})$. Finally note that $Supp(\nu(\cdot \mid \mathbf{x})) = \mathbf{D}$ for all $\mathbf{x}$, which proves (iii). Consequently Theorem 6.3.3 applies for $\{(\chi^N(t))_{t\in[0,1]}\}_N$.

Let now ρ be defined by

$$\rho(\mathbf{x}, \mathbf{y}) := \inf\left\{\mathbf{I}(\mathbf{x}, \varphi) : \ \varphi \in \mathcal{C}([0,1], \mathbb{R}^k), \ \varphi(1) = \mathbf{y}\right\}, \tag{24}$$

By a direct application of the *contraction principle* (see e.g. Theorem 4.2.1 in Zeitouni and Dembo (1998)), $\mathbf{y} \mapsto \rho(\mathbf{x}, \mathbf{y})$ has compact level sets. Moreover, the sequence $(\chi^N(1))_N$, with $\chi^N(0) = \mathbf{x}$, satisfies a large deviation principle with rate function $\rho(\mathbf{x}, \cdot)$:

- for any open set $\mathbf{U} \subseteq \mathbb{R}^k$,

$$\liminf_{N \to +\infty} \frac{1}{N} \log \mathbb{P}\left(\chi^N(1) \in \mathbf{U} \mid \chi^N(0) = \mathbf{x}\right) \geq - \inf_{\mathbf{y} \in \mathbf{U}} \rho(\mathbf{x}, \mathbf{y}),$$

- for any closed set $\mathbf{C} \subseteq \mathbb{R}^k$

$$\limsup_{N \to +\infty} \frac{1}{N} \log \mathbb{P}\left(\chi^N(1) \in \mathbf{C} \mid \chi^N(0) = \mathbf{x}\right) \leq - \inf_{\mathbf{y} \in \mathbf{C}} \rho\left(\mathbf{x}, \mathbf{y}\right),$$

  uniformly for $\mathbf{x} \in K$.

Also, if $\rho(\mathbf{x}, \mathbf{y}) = 0$ then there exists $\varphi \in \mathcal{C}([0,1], \mathbb{R}^k)$ such that $\varphi(1) = \mathbf{y}$ and $I(\mathbf{x}, \varphi) = 0$. Hence, by definition of $\mathbf{I}$, we necessarily have $\varphi(t) = \psi(\mathbf{x}, t)$, for $t \in [0,1]$. Thus $\mathbf{y} = \varphi(1) = \psi(\mathbf{x}, 1) = F(\mathbf{x})$. Reversely, if $\mathbf{y} = \psi(\mathbf{x}, 1)$ then $\rho(\mathbf{x}, \mathbf{y}) \leq \mathbf{I}(\mathbf{x}, \psi(\mathbf{x}, \cdot)) = 0$. Finally $\rho(\mathbf{x}, \mathbf{y}) = 0$ if and only if $\mathbf{y} = \psi(\mathbf{x}, 1)$.

Since kernels μ and ν coincide on $\mathbf{S}_*$, the family $\left(\mathbf{x}_N^N\right)_{N \geq 1}$ coincides with $\left(\chi_N^N\right)_{N \geq 1}$ (that is, $(\chi^N(1))_N$), as long as it remains in $\mathbf{S}_*$. Therefore, if $\mathbf{x} \in \mathbf{S}$ then $\rho(\mathbf{x}, \mathbf{y}) = 0$ iif $\mathbf{y} = \psi(\mathbf{x}, 1) = \phi(\mathbf{x}, 1) = F(\mathbf{x})$. Moreover, $\left\{\mathbf{x}_N^N\right\}_{N \geq 1}$ satisfies a uniform large deviation principle with rate function ρ, on $\mathbf{S}_*$:

- for any open set $\mathcal{U} \subseteq \mathbf{C}([0,1], \mathbf{S}_*)$,

$$\liminf_{N \to +\infty} \frac{1}{N} \log \mathbb{P}\left(\mathbf{x}_N^N \in \mathbf{U} \mid \chi^N(0) = \mathbf{x}\right) \geq - \inf_{\varphi \in \mathbf{U}} \rho\left(\mathbf{x}, \varphi\right),$$

- for any closed set $\mathcal{C} \subseteq \mathbf{C}([0,1], \mathbf{S}_*)$

$$\limsup_{N \to +\infty} \frac{1}{N} \log \mathbb{P}\left(\mathbf{x}_N^N \in \mathbf{C} \mid \chi^N(0) = \mathbf{x}\right) \leq - \inf_{\varphi \in \mathbf{C}} \rho\left(\mathbf{x}, \varphi\right),$$

  uniformly for $\mathbf{x} \in K$.

This conclues the proof that Assumption 3 holds. ■

Remark 2 *Notice that when the kernel $\mu(\cdot \mid \mathbf{x})$ depends on the parameter N, we need to use the more general results of chapter 6.7 in Dupuis and Ellis (2011), where they extend Theorem 6.3.3 to models with a residual noise term. However we avoid unnecessary technicalities by assuming that it does not depend on N.*

X.3 Other proofs

Lemma 1 *Assumption 2 holds with $V_c = [0, \gamma(c)[\times[0, \gamma(c)[\cup]1 - \gamma(c), 1]\times]1 - \gamma(c), 1]$, for some $\gamma(c) > 0$.*

Proof of Lemma 1: Let $c > 0$ and define $U_c := \{\mathbf{x} \in \mathbf{S} : \ x_i < \gamma(c), \ \forall i = 1, ..., k\}$, where the choice of $\gamma(c)$ will be made clear in a moment. Assume also without loss of generality that $p_i(\mathbf{1} - \mathbf{x}) \geq 1/2$ for any i and any $\mathbf{x} \in U_c$. Let $\bar{\mathbf{x}} := (n/N, n/N)$, with $n = \gamma(c)N$. Then

$$\mathbb{P}(\mathbf{x}_N^N = \mathbf{0} \mid \mathbf{x}_0^N = \bar{\mathbf{x}}) \geq \mathbb{P}\left(\mathbf{x}_1^N = \mathbf{x}^{N,1}, ..., \mathbf{x}_n^N = \mathbf{x}^{N,n} \mid \mathbf{x}_0^N = \bar{\mathbf{x}}\right),$$

where $\mathbf{x}^{N,j} = ((n-j)/N, (n-j)/N)$. As a consequence,

$$\begin{aligned}
\mathbb{P}(\mathbf{x}_N^N \in \mathbf{S}_0^N \mid \mathbf{x}_0^N = \bar{\mathbf{x}}) &\geq \prod_{j=0}^{n-1} \left(\frac{(n-j)^2}{N^2} p_1(\mathbf{1} - \mathbf{x}^{N,j}) p_2(\mathbf{1} - \mathbf{x}^{N,j}) \right) \\
&\geq \frac{(n!)^2}{(2N)^{2n}} \\
&\geq n \left(\frac{n}{2eN}\right)^{2n} \\
&\geq \gamma(c) N \frac{\gamma(c)}{2e}^{2\gamma(c)N}
\end{aligned}$$

because $n! \geq n^{n+1/2}e^{-n}$ for any $n \in \mathbb{N}$. Let us choose $\gamma(c)$ small enough so that $k\gamma(c)\log\gamma(c) - k\gamma(c)(1+\log 2) \geq -c$. We then have

$$\begin{aligned}\frac{1}{N}\inf_{\mathbf{x}\in U_c}\log\mathbb{P}(\mathbf{x}_N^N \in \mathbf{S}_0^N \mid \mathbf{x}_0^N = \mathbf{x}) &= \frac{1}{N}\log\mathbb{P}(\mathbf{x}_N^N \in \mathbf{S}_0^N \mid \mathbf{x}_0^N = \bar{\mathbf{x}}) \\ &\geq \frac{1}{N}\log\gamma(c) + 2\gamma(c)\log\left(\frac{\gamma(c)}{2e}\right)\end{aligned}$$

which gives the desired result, by taking $\gamma(c)$ small enough so that $2\gamma(c)\log\left(\frac{\gamma(c)}{2e}\right) \geq -c/2$. ■

Proof of Theorem 3. The proof simply follows by applying Theorem 1 and 2. ■

Lemma 2 *Assumption 2 holds, with* $V_c = [0, \underline{R} + \gamma(c)[$*, for some* $\gamma(c) > 0$.

Proof of Lemma 2. Let $c > 0$ and $V_c := [0, \underline{R} + \gamma(c)]$, where $\gamma(c) < \frac{-c}{\log(\underline{R}(1-\pi(\underline{R}))}$. If $R_0^N = R \in]\underline{R}, \underline{R} + \gamma(c)[$, sufficient condition to have $R_N^N \leq \underline{R}$ is that the first $[N\gamma(c)] + 1$ reviews are bad. Since $R(1-\pi(R)) \geq \underline{R}(1-\pi(\underline{R}), \forall R \geq \underline{R}$,

$$\frac{1}{N}\log\mathbb{P}(R_N^N \in \mathbf{S}_0^N \mid R_0^N = R) \geq \frac{[N\gamma(c)+1]}{N}\log\left(\underline{R}(1-\pi(\underline{R})\right).$$

For N large enough, $\frac{[N\gamma(c)+1]}{N} \leq \frac{-c}{\log(\underline{R}(1-\pi(\underline{R}))}$. It concludes the proof. ■

Proof of Theorem 4. Since the map $R \mapsto \pi(R) - R$ is strictly decreasing on $]\underline{R}, 1]$, we have

$$\pi(\underline{R}) := \lim_{R\to\underline{R}^+}\pi(R) = \mathbb{P}\left[\rho(\overline{\theta}, \zeta_n, Q) \geq 0\right],$$

since, by continuity of ρ, when R goes to $\underline{R}$ from above, only customers of type $\overline{\theta}$ buy the product.

- If $Q = 1$ then $\pi(\underline{R}) > \underline{R}$. Moreover, $\pi(1) \leq 1$, and the unique solution of $\pi(R) = R$ is an attractor for the system $\dot{R} = -R + \pi(R)$.
- If $Q = 0$ then $\pi(\underline{R}) < \underline{R}$, hence $\pi(R) - R < 0$, $\forall R \geq \underline{R}$ and the process will drive R below $\underline{R}$. ■

Proof of Theorem 5. This is a direct consequence of both Theorem 1 and 2. ■

XI Appendix B: Microfoundations of the Opinion Formation Model

In this application, we consider the following: at period n, an individual i holds opinion $y_i(n) \in \{0, 1\}$, where $y_i(n) = 0$ means that i currently supports candidate B, and $y_i(n) = 1$ means that i currently supports candidate A. When individual i is drawn at random, he will interact with a (randomly chosen) set K of individuals, with $|K| = k$.

Individuals have preferences both for consistency, meaning that it is costly for them to change opinions, and for conformity, meaning that they dislike holding an opinion which is very different from what they observe in society. We model this as follows: individual i has to choose between $y_i(n+1) = 0$ or $y_i(n+1) = 1$, by maximizing the quantity:

$$y_i \in \{0,1\} \mapsto -\left[y_i - \frac{1}{k+\alpha_i}\left(\alpha_i y_i(n) + \sum_{j \in K} y_j(n)\right)\right]^2$$

where α_i is the weight that individual i assigns to his own current opinion in his choice of his future opinion (the consistency parameter). If $K = \emptyset$, i.e. there are no interactions, then individual i would always choose $y_i(n+1) = y_i(n)$, by consistency. When $K \neq \emptyset$, the higher α_i, the harder for individual i to change opinion.

The second term in brackets is a weighted sum of current opinions in the sample drawn by individual i including his weighted self. Individual i wants to minimize the difference between his opinion and this weighted average.

Direct computations show that an individual i will change opinion after social interactions if and only if a sufficiently large share of the individuals he interacts with share the opposite opinion. Specifically,

$$y_i(n) = 0 \text{ and } y_i(n+1) = 1 \iff k + \alpha_i < 2\sum_{j \in K} y_j(n)$$

Similarly,

$$y_i(n) = 1 \text{ and } y_i(n+1) = 0 \iff k + \alpha_i > 2(\alpha_i + \sum_{j \in K} y_j(n))$$

So, when more than half the individuals, i included, share the opposite opinion from i, he will choose to change his opinion. In any other circumstances he would choose to stick to his current opinion. For instance, if $k = 6$ and $\alpha_i = 3$, then the individual interacts with 6 other individuals and considers himself as having a weight equivalent to 3 other individuals. In that case, a B supporter will change to A if and only if 5 or 6 out of the 6 people he interacts with are A supporters. In any other case he would stick to supporting B.

The case we focus on in the main text is for $k = 2$ and $\alpha_i = 1$, meaning that individuals interact with 2 other individuals and consider themselves as having the same weight as any other individual. In that case, a B supporter will change to A if and only if $\sum_{j \in K} y_j(n) = 2$, i.e. the two individuals he meets are A supporters.

With x_i denoting the share of A supporters in community i, the probability with which a B supporter will meet two A supporters out of two interactions is thus x_i^2, leading to equation (10) in the main text.