Dangerous tangents: an application of Γ-convergence to the control of dynamical systems

Rosario Maggistro* Paolo Pellizzari† Elena Sartori‡
Marco Tolotti§

Abstract

Inspired by the classical riot model proposed by Granovetter in 1978, we consider a parametric stochastic dynamical system describing the collective behavior of a large population of interacting agents. By controlling a parameter, a policy maker aims at maximizing her own utility which, in turn, depends on the steady state of the system. We show that this economically sensible optimization is ill-posed and illustrate a novel way to tackle this practical and formal issue. Our approach is based on Γ-convergence of a sequence of mean-regularized instances of the original problem. The corresponding maximizers converge towards a unique value which intuitively is the solution of the original ill-posed problem. Notably, to the best of our knowledge, this is one of the first applications of Γ-convergence in economics.

Keywords: Dynamical Systems, Finite Population Dynamics, Γ-convergence, Saddle-node Bifurcations, Social Interaction

JEL Classification: C61, C63, D91

MSC Classification: 91B14; 49J45

*University of Trieste, Department of Economics, Business, Mathematics and Statistics, Via dell'Università 1, 34127 Trieste, Italy. Contact: rosario.maggistro@deams.units.it.

†Ca' Foscari University of Venice, Department of Economics, Cannaregio 873, 31121, Venice, Italy. Contact: paolop@unive.it.

‡University of Padova, Department of Mathematics "Tullio Levi-Civita", 63, Via Trieste, 35121, Padova, Italy. Contact: esartori@math.unipd.it.

§Ca' Foscari University of Venice, Department of Management, Cannaregio 873, 31121, Venice, Italy. Contact: tolotti@unive.it.

1 Introduction

Understanding and controlling collective behavior is a challenging, subtle and potentially very useful endeavour (see Granovetter [1978]). The mere attempt to know agents' preferences, motives and norms is fraught with practical and conceptual difficulties: data are sometimes scarce or difficult to obtain and persons may have blurred incentives and hide or misrepresent, consciously or not, the drivers of their behavior. Moreover, even when a reasonable model for individual actions is agreed, results can be puzzling as shown, for instance, in Thomas Schelling's segregation model where *slight* homophilic preferences for neighbors of the same race lead in equilibrium to *massive* residential segregation, Schelling [1971]. These collective models are know to produce very diverse and intriguing outcomes (simple vs complex attractors, bifurcations, stable polarized vs non-polarized equilibria,...). Therefore, a policy maker (or any external mastermind) could be interested in *controlling* the behavior of the population by setting, in her capacity, some parameters at some *optimal level*.

We consider a problem characterized by two components: a dynamical system describing the evolution of some state variable that represents, in aggregate terms, agents' collective behavior; and a social planner who *ex-ante* sets the parameters of such dynamics to maximize her payoff which, in turn, depends on the steady state of the system. Interestingly, in some cases, the problem turns out to be *ill-defined* in that the objective function is bounded form above but does not admit a maximum. However, we show that, by defining a proper sequence of auxiliary stochastic problems, it is possible to formally prove a convergence result whose limit is well-defined and helps in shedding light on the original problem.

To fix ideas and keeping for the moment the formalism at a minimum, consider the dynamical system

$$r(t+1) = F(r(t); \sigma),$$

where the future state $r(t+1)$ depends on the current state $r(t)$ through the (iterated) application of F, as well as on a parameter σ. Assume now that for a given initial condition $r(0) = r_0$, a limit state $\rho(\sigma)$ is reached.[1] The social planner attempts to solve an optimization problem which depends on $\rho(\sigma)$ and is subject to some costs $c(\sigma)$. Specifically:

$$\max_{\sigma} \quad f(\sigma), \text{ where } f(\sigma) = \rho(\sigma) - c(\sigma). \tag{1}$$

[1]Formally, this means that $\rho(\sigma) = \lim_{t\to\infty} F^t(r_0; \sigma)$, where F^t is the composition of F with itself for t times.

The fact that the payoff of the social planner positively depends on the long-run outcome of the system, (i.e., on $\rho(\sigma)$) seems reasonable for many applications. Consider, for example, the case where $r(t)$ represents the market share of a durable good and σ is a parameter related to the quality of the same good. In this respect, $\rho(\sigma)$ can be interpreted as the equilibrium market share related to a specific quality strategy implemented by the company. In the payoff, the cost component, $c(\sigma)$, plays the role of the (potential) cost to implement such strategy and, mathematically speaking, completes the definition of a specific parametric dynamical system indexed by σ.

Modelling approaches of individual decisions giving rise, at the aggregate level, to dynamics resembling (1) can be found, among others, in Blume & Durlauf [2003] or Barucci & Tolotti [2012]. Models of supply and demand emerging in financial markets, where a similar structure manifests itself are presented in Gordon et al. [2013]. Diffusion of innovations and adoption models of durable goods as pioneered in Bass [1969] and, recently, in Peres et al. [2010] share a micro-foundation which is similar in spirit to the one presented here. Also epidemiological models of SI (Susceptible Infected) type exhibit dynamics which resemble the formalism proposed in (1). Finally, as already said, a celebrated model in this context goes back to Granovetter [1978]; here, a mastermind *manipulates* the mood of crowds to trigger riots that in the end involve a large fraction of the population. Hereafter, due to its clarity and simplicity, this latter riot model will be used to motivate our treatment and exemplify our results.

The previous situations all share a similar conceptual structure: contingent on the value of the parameter, some equilibrium is likely to endogenously appear as the final outcome of the dynamics and an external agent is interested in shaping or controlling this outcome. While the intuition appears to be quite natural, the explicit setup of this "optimization problem over a dynamical system" is not so frequent and we explore some of the intrinsic difficulties that may arise defining such problems.

In fact, we show that problems formalized as in (1) can be *ill-posed* in that the objective function is discontinuous, has no maximum but only admits a supremum. Essentially, this is due to a *saddle-node bifurcation* characterizing the dynamical system: one of the fixed points of the equation $r = F(r; \sigma)$ disappears, when σ increases beyond some critical value σ_c.[2] Note that $\rho(\sigma)$ is one of such fixed points (actually, the smallest one in our setup) and the presence of a bifurcation makes its value abruptly jump at the bifurcation value. On the top of that, as already noticed in Granovetter [1978], the

[2]A brief description of saddle-node bifurcations is postponed to Appendix A. See Strogatz [2015] for an exhaustive analysis of bifurcations in dynamical systems.

bifurcation value σ_c is intuitively the value the value the masterind is looking for. In loose words, the system is purposedly steered to reach a bifurcation point (as this is advantageous).

However, as a matter of fact, this value is not a maximizer for f, but rather a point where a supremum is attained. As a consequence, this optimization problem is intrinsically ill-posed. Summarizing, at odds with their natural appeal, problems like (1) can be fraught with technical difficulties that may impede formal and numerical treatments.

Technically speaking, one of the main goals of this paper is to overcome such technical issue by examining the problem from a different angle. Instead of looking at the steady states of the one-dimensional (deterministic) dynamical system as expressed in (1), we introduce a regularized stochastic version of it, where the number of agents in the population is large but finite. Albeit stochastic, this problem is now well-defined and admits a unique maximizer σ_N^* for all N, where N is the size of the population. Finally, we are able to provide a formal limit for N going to infinity and to show that the sequence of maximizers, $(\sigma_N^*)_{N\geq 1}$, converges exactly to σ_c, the value that is "expected" to be a solution of the ill-posed original problem. Interestingly, it turns out that the aforementioned convergence holds in a Γ-*sense*: we will formally prove that the sequence of objective functions f_N of the auxiliary N-dimensional problems Γ-converges to a well-defined f_∞ (see (13)). This Γ-limit turns out to be *almost everywhere* equal to the original objective function f except at the optimal level σ_c.[3] The result sheds light on one rigorous way to deal with the ill-posedness of the original problem hinted at in Granovetter's work. More recently, singularities similar to the one described above are also detected in Nadal et al. [2005] and Gordon et al. [2013]. In these works, the authors model a monopolist in charge to set the optimal price to foster demand and maximize profits. With language and notation borrowed often from physics, these papers contain a model that has similarities with ours and contains a lengthy discussion of the "*epistemic uncertainty*" inherent in selecting one price to maximize monopolist's profits. Uncertainty due to the presence/disappearance of multiple equilibria is acknowledged by sentences like "*[the optimal solution] lies very near the critical price value at which such high demand no more exists"* (Gordon et al. [2013]). The singularity recognized (albeit not "solved") in their model can be formally tackled through the use of a Γ-convergent sequence of problems whose limit results to be almost everywhere identical to the original singular model.

Coming to the methodology, to the best of our knowledge, papers in the

[3] We refer the reader to Appendix B for a brief recap on Γ-convergence and to Braides [2002] for a complete treatment.

realm of mathematical economics and social sciences employing Γ-convergence are quite rare. Ghisi & Gobbino [2005] describes a variational problem arising from a generalization of the well known monopolist's problem introduced in Rochet & Choné [1998]. In this model, the monopolist proposes a set of products and looks for the optimal price list which minimizes costs, hence maximizing the profit. This leads to a minimum problem for functionals H (the "pessimistic cost expectation") and G (the "optimistic cost expectation"), which are in turn defined through two nested variational problems. The authors prove that the minimum of G exists using an approximating sequence which $\Gamma-$converges to G, and that such a minimum coincides with the infimum of H. An economic model of monopoly has also been studied in a general setting by Monteiro & Page Jr. [1998], and under convexity assumptions by Carlier [2001]. In the latter, the author studies a principal-agent model with adverse selection and characterizes the incentive-compatible contracts in terms of an envelope property called $h-$convexity. Using this characterization, the principal's problem is written as a non-standard variational problem (with $h-$convexity constraint) for which existence of a solution is proved. In Monteiro & Page Jr. [1998], similar general existence results for the principal-agent problem with adverse selection are given. However, the way to deal with the problem is different from the one in Carlier [2001], since the authors consider budget constraints that force the prices to remain in a given compact set and their results rely on a *nonessentiality assumption* (i.e., nonessentiality of some goods relative to other ones) which is not required in Carlier [2001].

We show in this paper how Γ-convergence can be used to deal with a natural but ill-posed problem faced by the decision-maker. Indeed, the economic interest of our treatment stems from the observation that it is the decision-maker herself who rationally pushes the system to configurations where some equilibria disappear. This fact also poses severe difficulties when numerical solutions are searched, especially if one studies higher dimensional systems, where the insightful graphical representations of Granovetter [1978] are unavailable. The decision-maker is akin to someone willing to reach the edge of a cliff to reap benefits while approaching, at the same time, a dangerous discontinuity.

The paper is organized as follows. Section 2 describes in detail the optimization problem as stated in (1). In Section 3 we provide a stochastic version of the same problem where the number of actors in the economy is now finite. This stochastic approach, albeit looking more cumbersome at the surface, provides a double benefit: first, it allows us to take advantage of probabilistic tools and, secondly, it naturally leads the modeler to simulation and numerical methods. We will see that this finite-dimensional approach is crucial to analytically set the proper convergence scheme. Section 4 is devoted to the analysis of numerical simulations, whereas in Section 5 we draw some conclusions. Appendix A contains all technical proofs and Appendix B summarizes the main concepts of Γ-convergence employed in the paper.

2 The deterministic dynamic model

Inspired by Granovetter [1978], we first consider a (infinite) population of actors facing the decision of taking part in a riot. Agents are heterogeneous, they are aware of the actual proportion of people involved in the riot and decide whether to join or not according to a personal (random) activation threshold. If the proportion of active agents is above the threshold, they join the crowd; otherwise they don't. Technically speaking, the random thresholds are independent copies of a random variable with distribution F, which is assumed to be, from now on, Gaussian with fixed mean $\mu = 0.25$ and standard deviation $\sigma \geq 0$ [4]. Recall that σ will be set by the social planner. In this respect, by choosing σ, she is fixing a specific Gaussian distribution F and, hence, a specific *behavioral trait* of the underlying population of potential rioters. Practically, this can be interpreted as a purposely inflation of the volatility of the crowd to ignite a riot.

Let call $r(t)$ the proportion of agents taking part in the riot by time t, where $t \in \mathbb{N}$ and fix $r(0) = 0$.[5] People, whose threshold is lower than $r(t)$, decide to join the riot at time $t+1$; moreover, once they are part of the riot, it is not possible to retire. Therefore, $r(t+1)$ will account for the people already involved at previous time plus the proportion of *newcomers*. In other terms, this is exactly $F(r(t); \sigma)$, i.e., the probability of the threshold being lower than $r(t)$. Hence, we obtain the recursive equation of Granovetter [1978]:

$$r(t+1) = F(r(t); \sigma). \tag{2}$$

We denote by $\rho(\sigma)$ the smallest solution[6] of the fixed point equation

$$r = F(r; \sigma). \tag{3}$$

[4]We use the statistical specification suggested in Granovetter [1978].

[5]We are interested in studying the case where the initial condition is zero (or small) as a benchmark for applications where the social phenomenon is analyzed from its beginning.

[6]In this respect, $\rho(\sigma)$ can be interpreted as the *lowest fixed point* function in the sense of Milgrom & Roberts [1994]. In that classical paper, the authors discuss conditions for monotonicity of such a map. Conversely, we are interested in those situations where the monotonicity is broken.

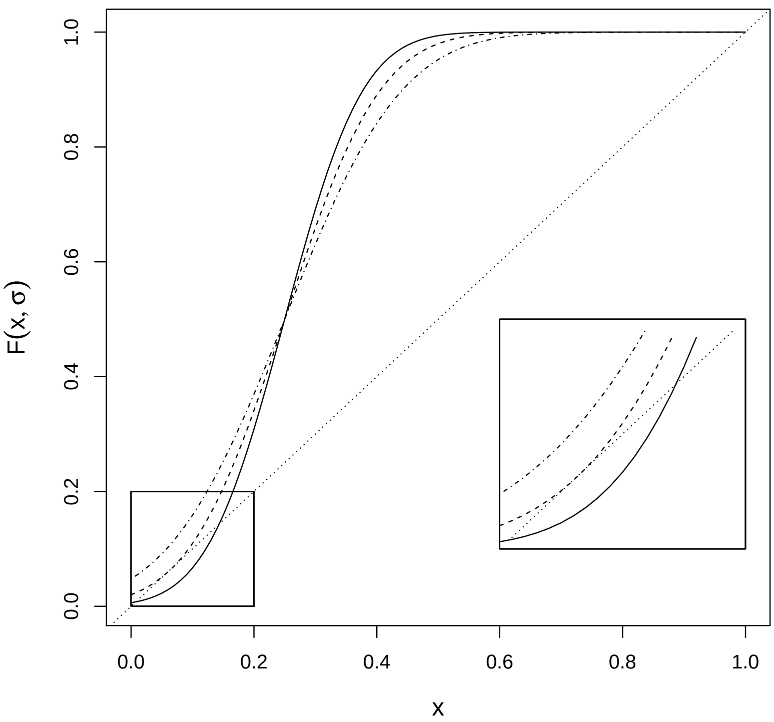

Figure 1: The graph of $x \mapsto F(x;\sigma)$ for three different values of σ. Specifically, $\sigma = 0.1$ (solid line), $\sigma = 0.122$ (dashed line), and $\sigma = 0.15$ (dot-dashed line).

In case F admits a unimodal density function (as in our example based on Gaussian thresholds), it is possible to characterize such an equilibrium. In fact, it is well known that under this minimal assumption, there are at most three solutions to (3), and the number of such solutions depends on the value of σ. Moreover, when there are three equilibria, the intermediate one is always unstable whereas the two extreme ones are (locally) stable. As said, we are interested in the case where the initial condition is zero and, therefore, the dynamical system always converges to the smallest solution of (3). This family of dynamical systems exhibits what is called a *saddle-node bifurcation*.[7] Keeping $\mu = 0.25$, the black line in Figure 1 shows three fixed points of $r = F(r;\sigma)$, occurring when $\sigma = 0.1$. The intermediate fixed point,

[7]For a formal analysis of saddle-node bifurcations and related facts, we refer the reader to Lemma A.1 in Appendix A.

located at about 0.15 is unstable and, hence, the limit state for any initial condition below this fixed point would converge to the lower fixed point. For all such initial conditions, the limiting fraction of population involved in the riot would be about 0. Suppose now that the social planner has the power to increase σ (at a cost): this has the effect to slide upwards the graph of F, up to the point where the two smaller fixed points merge when F is tangent to the bisector line (the case is depicted with the dashed line in Figure 1, obtained when $\sigma = 0.122 \approx \sigma_c$, where σ_c denotes the exact tangency value). Now, the limit state when the initial fraction of citizens involved in the riot is an infinitesimal fraction would raise to about 6.2%. But, even more importantly, a further tilt to F beyond σ_c triggers the occurrence of a saddle-node bifurcation and the limit state abruptly jump to 1, see the dot-dashed curve in the picture, relative to $\sigma = 0.15$.

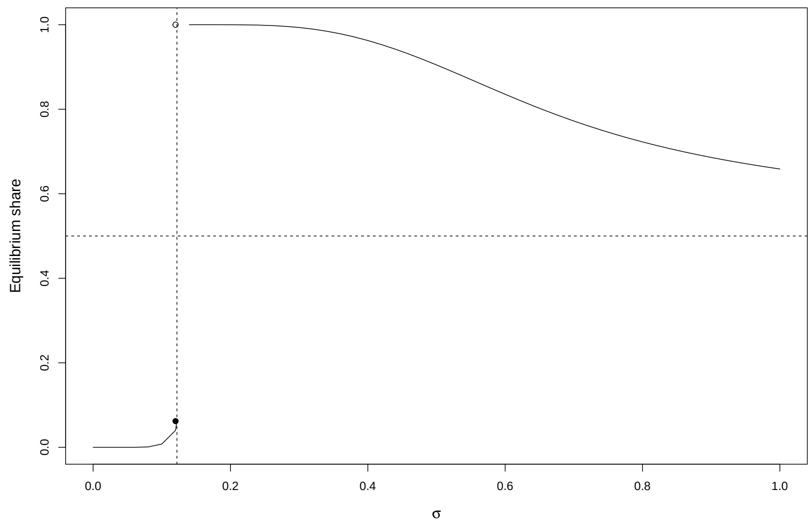

Figure 2: Equilibrium share $\rho(\sigma)$ of the population involved in the riot as a function of σ. The discontinuity occurs at $\sigma_c \approx 0.122$.

Figure 2 depicts the effect of the saddle-node bifurcation on the equilibrium share $\rho(\sigma)$ of the model: clearly, as σ is increased beyond σ_c, the disappearance of two fixed points generates a jump in the limit state that will be reached through iteration of F. As pointed out in Strogatz [2015], the term "saddle-node" is not entirely consolidated in the dynamical systems' literature and, for instance, such a bifurcation is quite imaginatively called a "blue sky bifurcation" in some cases to stress that, reversing the direction of the change of σ, a new equilibrium can be created "out of the blue".

Figure 2 shows that $\rho(\sigma)$ has no maximizer as a function of σ, that can be interpreted as the *mood of the population*. As it was influentially pointed out in Granovetter [1978], finite populations drawn from the same F, when $\sigma \approx \sigma_c$, may lead to very different limit states under iteration: "*There is no obvious sociological way to explain why a slight perturbation of the normal distribution around the critical standard deviation should have a wholly discontinuous, striking qualitative effect... This example shows again how two crowds whose average preferences are nearly identical could generate entirely different results*".[8]

The revelation that minor perturbations in individual features could produce large aggregate effects was probably a part of the *zeitgeist* of the late seventies as similar ideas are present in Schelling [1971] concerning racial segregation. Likewise, Allen and Sanglier in 1979, examining dynamic models of urban growth distilled their results observing that "*small perturbations of density, perhaps of random origin, are amplified by the interactions between the elements of the system, and lead to a qualitative change in the macroscopic structure of the spatial distribution*".[9] It is opinion of the authors that these insights keep their freshness despite four decades have passed.

To make this argument formal, we deviate from the original paper and introduce an *external agent* acting as a social planner. Suppose that a *rabble-rouser* is interested in fueling a massive participation in the uprising. Of course, due to the cost of stiring-up people, the policy maker is facing the following decision problem: *"How to optimally settle the mood of the population in order to obtain a large-scale riot?"*. Consider a unitary cost $k > 0$ for a unit of standard deviation, her optimization problem can be written as

$$\max_{\sigma} f(\sigma), \text{ where } f(\sigma) = \rho(\sigma) - k\sigma. \qquad \text{(P1)}$$

Notice that, if k is too large, the problem loses of interest since the objective function reaches a proper maximum at some point $\sigma \leq \sigma_c$, possibly, at $\sigma = 0$. Therefore, in the following lemma we define a suitable threshold for k. A similar argument applies also to the expected value μ related to the distribution function F.[10] Although no closed-form solution for this problem exists, the objective function has the same behavior depicted in Figure 2, as proven by the following lemma.[11]

---

[8] See Granovetter [1978], pages 1427–1428.

[9] See Allen & Sanglier [1979], page 257.

[10] More details on such threshold levels for k and μ are provided in Appendix A.

[11] All proofs are postponed to Appendix A.

Lemma 2.1. *Define σ_c as the value of σ that makes the graph of the map $x \mapsto F(x;\sigma)$ tangent to the bisector line. Define also*

$$\tilde{f} = \lim_{\sigma \to \sigma_c^+} f(\sigma) = \tilde{r}_c - k\sigma_c, \quad \tilde{r}_c := \lim_{\sigma \to \sigma_c^+} \rho(\sigma). \tag{4}$$

Suppose finally that $k < k^{th} := \frac{\tilde{r}_c - \rho(\sigma_c)}{\sigma_c}$ and $\mu < \frac{1}{2}$. Then, the function

$$f(\sigma) = \rho(\sigma) - k\sigma,$$

where $\rho(\sigma)$ is the smallest solution to (3), has the following properties:

i) it is left-continuous with a discontinuity point at σ_c;

ii) it admits a finite supremum $\tilde{f}$.

In other words, f has no maximum and the sup is reached, for $\mu = 0.25$, at $\sigma_c \approx 0.122$. On the other hand, σ_c is exactly the value, where the curve $F(x;\sigma_c)$ is tangent to the bisector line (see Figure 1). Algebraically, σ_c is the unique value of the parameter such that the following system is solvable for some $r \in [0,1]$:

$$\begin{cases} F(r,\sigma) = r, \\ \dfrac{dF}{dr}(r;\sigma) = 1. \end{cases}$$

As a matter of fact, the problem is clearly ill-defined: no solution exists, even though the "optimal point" has a very clear geometrical interpretation in terms of the distribution function of the random thresholds.

To deal with this issue, we propose in the next section an auxiliary optimization problem. In particular, we let the population of actors to be finite of size N. We will see that this makes the problem stochastic but, on the other hand, well-posed for any finite N. Eventually, we will show that the sequence of objective functions of such finite dimensional problems converges in a Γ-sense to a function f_∞ which will be almost everywhere equal to the original objective function f. However, differently from f, f_∞ admits a unique maximum; therefore, the new auxiliary problem will be well-defined also for $N \to \infty$.

3 The (stochastic) problem with N agents

Consider N actors facing the following decision problem: "to participate or not into a *riot*". The state variable is thus binary. We define $y_i(t) \in \{0,1\}$ for $i = 1,\dots,N$, where $y_i(t) = 1$ means that agent i is participating at the

riot at time t. Once part of the riot, it is not possible to retire. Moreover, we assume that $y_i(0) = 0$ for all i. This choice captures (makes manifest) the fact that, in the original infinite-dimensional problem, we analyze the social phenomenon from its beginning when, with the position $r(0) = 0$, it is assumed that no agent is initially rioting.
Being a decision under *collective behavior*, the reward in participating at the riot depends on the present number (or better, the proportion) of people involved. This quantity is given by[12]

$$r_N(t) = \frac{\sum_i y_i(t)}{N}. \tag{5}$$

Any actor decides to join as soon as the quantity r_N is large enough. We model *random thresholds* X_i, $i = 1, \ldots, N$ as N independent copies of a random variable X with absolutely continuous distribution function F. The rule is straightforward:

$$y_i(t+1) = 1 \quad \Leftrightarrow \quad X_i \leq r_N(t). \tag{6}$$

Summing up and normalizing over N, one sees that

$$r_N(t+1) = \frac{1}{N}\sum_{i=1}^{N} y_i(t+1) = \frac{1}{N}\sum_{i=1}^{N} \mathbb{I}_{\left\{(-\infty, r_N(t)]\right\}}(X_i). \tag{7}$$

The right hand side of equation (7) is nothing but $F_N(r_N(t);\sigma)$, the N-dimensional empirical distribution of the random thresholds. We thus obtain the following recursive equation, characterizing the N-dimensional system

$$r_N(t+1) = F_N(r_N(t);\sigma). \tag{8}$$

We now have a population of N agents evolving according to (8). On the top of that, the social planner is aiming at controlling the behavior of the population by setting the *optimal* σ so to maximize a measurable function of the stochastic process r_N. More specifically, the policy maker considers the *steady states* of such dynamics as a natural outcome to consider.

Should the social planner look at the *finite-dimensional* (stochastic) population or, rather, implement some theoretical asymptotic results to rely on some simpler deterministic limit dynamics? Technically speaking, this translates in the order in which we calculate the limits of $N \to \infty$ and $t \to \infty$. To

[12]We could also consider the quantity $r_N^{(i)} = \frac{\sum_{j\neq i} y_j}{N-1}$. When N becomes large (infinite) the contribution of y_i is negligible, thus the two problems have exactly the same limiting behavior.

deal with a deterministic system, we first take the N-limit; the computation of steady states is then straightforward but, as just seen above, the optimization problem is ill-posed. Conversely, if we first perform the time-limit, the observable steady states are stochastic and the modeler relies on statistical tools to remove the randomness inherent in the finite-dimensional system. In the remainder of the section, we formalize the two approaches, stressing how different they are in their outcomes.

3.1 Weak convergence of stochastic processes

The first approach relies on the classical asymptotic theory developed in Ethier & Kurtz [2009]. The convergence of the stochastic process r_N to a well-posed limiting process r, when the number of actors N tends to infinity, is ensured by the following result.[13]

Proposition 3.1. *Assume the recursion given in* (8), *with* $r_N(0) = 0$ *for all* N. *When* $N \to \infty$, *the process* $(r_N(t))_{t\geq 0}$ *weakly converges towards* $(r(t))_{t\geq 0}$, *characterized by the following asymptotic recursion*

$$r(t+1) = F(r(t); \sigma), \qquad r(0) = 0, \tag{9}$$

where F *is the distribution of* X_i, $i = 1, \ldots, N$ *as defined in* (6).

Since, by assumption, $r(0) = 0$, necessarily the dynamical system converges to $\rho(\sigma)$, i.e., the smallest among the (possibly multiple) solutions of

$$r = F(r; \sigma), \tag{10}$$

obtained considered the limit for $t \to \infty$ of (9). In particular, as discussed above, there is no solution to the maximization problem of the social planner.

Example 3.2 (The Granovetter setting). *Assume* $X \sim \mathcal{N}(\mu, \sigma)$, *where* μ *is fixed at the level* 0.25 *and* $\sigma > 0$. *In this case, there exists a critical level* σ_c *such that, for* $\sigma > \sigma_c$ *there is only one solution of* (3). *For* $\sigma < \sigma_c$, *the solutions are three and for* $\sigma = \sigma_c$ *the function* $r \to F(r; \sigma)$ *is tangent to the bisector and the solutions are exactly two.*

3.2 Γ-convergence of regularized operators

The second approach is based on the stochastic version of the dynamical system given in (8). Differently from what was done previously, we first take the

[13] The proof of this rather classical result is omitted. We refer the reader to Ethier & Kurtz [2009] for more details.

time-limit of the dynamics. As a matter of fact, the steady state is now a random variable, measurable w.r.t. the N-dimensional sample $X_1, \ldots, X_N$. Let the random variable R_N represent the (random) steady state of the dynamics in (8). Formally,

$$R_N = \lim_{t\to\infty} r_N(t). \tag{11}$$

To properly define an optimization problem for any finite N, we rely on the expected value of R_N as the observable to be investigated by the social planner:

$$\rho_N(\sigma) = \mathbb{E}^\sigma[R_N]. \tag{12}$$

It turns out that the auxiliary problem based on such new observable is now *well-posed*. Moreover, it is possible to characterize the limit of the sequence $(\rho_N(\sigma))_N$ in terms of $\rho(\sigma)$, the steady state related to the original and ill-posed optimization problem. These claims are stated in the following

Theorem 3.3. *Fix N and consider ρ_N as defined in* (12)*. Then, the optimization problem*

$$\max_\sigma f_N(\sigma), \qquad f_N(\sigma) := \rho_N(\sigma) - k\sigma, \tag{P2}$$

is well-posed and admits a maximizer σ_N^*. Moreover, for $N \to \infty$,

$$f_N \xrightarrow{\Gamma} f_\infty,$$

where

$$f_\infty(\sigma) = \begin{cases} \rho(\sigma) - k\sigma & \text{if} \quad \sigma \neq \sigma_c \\ \tilde{r}_c - k\sigma_c & \text{if} \quad \sigma = \sigma_c \end{cases} \tag{13}$$

and $\tilde{r}_c$ is defined in (4)*. As a consequence, the sequence of maximizers $(\sigma_N^*)_N$ related to (P2) converges to σ_c and $(\rho_N(\sigma_N^*))_N$ converges to $\tilde{r}_c$.*

The limit of f_N is thus taken in a Γ-sense and f_∞ is exactly the Γ-limit of the sequence of real functions $(f_N)_N$. It can be proven, instead, that the convergence is not properly working in a classical *almost sure* or *uniform* sense. Note, finally, that f_∞ is nothing but the upper-semicontinuous envelope of f (the original objective function of problem (P1)); the two functions only differ among each others for the value at σ_c which turns out to be exactly the maximizer for the problem where f_∞ is now used as a (auxiliary) objective function.

A summary of the different convergence schemes is graphically visualized in Figure 3. The starting point is the top-left stochastic dynamical system expressed by (8). Under the first approach, described in Section 3.1, we first move to the right by taking a N-limit, in a *"convergence of stochastic processes"* sense; then, we move down (time-limit), obtaining $\rho(\sigma)$. Moving in the area of the social planner, we come across the *ill-posed* problem (P1) as depicted in the bottom part of the right column. Conversely, by starting at the top-left corner and implementing the second approach described in Section 3.2, we first move down along the left column (time-limit), and later introducing the expectation operator, we arrive at the bottom-left point where σ_N^* is defined as the maximizer of f_N for any fixed N. Finally, as f_N converges in a Γ-sense to f_∞, we have that σ_N^* converges to σ_c, the maximizer of f_∞.

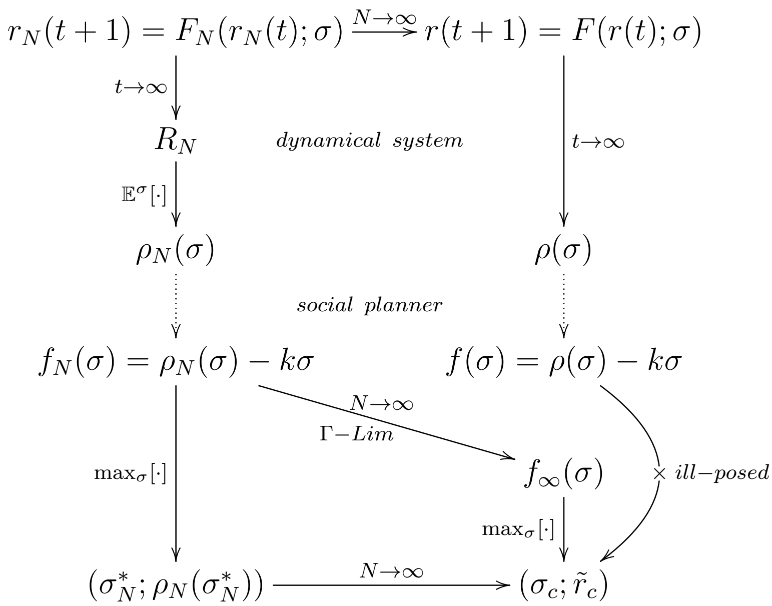

Figure 3: The scheme of convergences in the model.

In the next section we collect some numerical findings and give the sense of the previous asymptotic result solving the problem for increasing values of N. Moreover, we show that numerical results are notably affected by the accuracy in computing an estimator for the expectation of the random variable R_N.

4 Numerical findings

In this section, we numerically analyse the riot problem where, for concreteness, we set $k = 0.5$. We simulate M stochastic instances of the optimization problem (P2) in Theorem 3.3 to obtain M couples maximizer-maximum

$({}_m\hat{\sigma}_N^*;\, {}_m\hat{f}_N^*), m = 1, ..., M$, for given N. Then, we increase N to visually and numerically show the Γ-convergence mentioned in Theorem 3.3.

The key insight of the previous section was that convergence can be achieved by taking the expectation, $\rho_N(\sigma)$, of the random steady state R_N defined in (11). As a closed-form computation of $\rho_N(\sigma)$ appears to be impossible, we approximate it with a sample mean of S independent generations of the steady states. The numerical results are obtained as follows:

Given N;

for $m = 1, ..., M$;

Provide an estimate of the solution to Problem (P2), i.e., evaluate $f_N(\sigma)$ as many times as it is needed to get the optimal solution.[14]

1. To estimate $f_N(\sigma)$ for fixed σ:
   1.1 for $s = 1, ..., S$;
      1.1.1 sample individual activation thresholds ${}_s\tilde{X}_N$ from $\mathcal{N}(\mu, \sigma)$;
      1.1.2 compute the random equilibrium ${}_sR_N$;
   1.2 compute the sample mean $\hat{\rho}_N(\sigma) = \frac{1}{S}\sum_{s=1}^{S} {}_sR_N$;
   1.3 return the approximated value $\hat{f}_N(\sigma) = \hat{\rho}_N(\sigma) - k\sigma$;
2. Varying σ, identify the maximum ${}_m\hat{f}_N^*$ and the maximizer ${}_m\hat{\sigma}_N^*$.

Output: a M-dimensional vector of pairs $({}_m\hat{\sigma}_N^*;\, {}_m\hat{f}_N^*), m = 1, ...M$.

Since the sample mean converges to the correct average for $S \to \infty$, the higher the number S of simulated steady states is, the more our approximation will be close to $\rho_N(\sigma) = \mathbb{E}^\sigma[R_N]$. Hence, it is important to realize that the numerical method to solve the optimization problem of the social planner relies on an additional parameter, S, due to the need to replace the mathematical expectation with a sample mean. This amounts to say that, whenever an evaluation of the objective function in (3.3) is needed, a set of N thresholds will be independently re-sampled (over and over) for S times, as seen in the above pseudo-code to compute $f_N(\sigma)$.

Figure 4 shows some representative graphs of the objective functions of the optimization problem (P2). In the top panel, for fixed $S = 100$, the functions $\hat{f}_N(\sigma)$ relative to $N = 100, 1000$ and 10000 are depicted. It can be seen that larger values of N lead to smoother behavior and, more importantly, produce more accurate approximations of the function f_∞ defined in Theorem 3.3.

[14] We rely on the R function `optimize`.

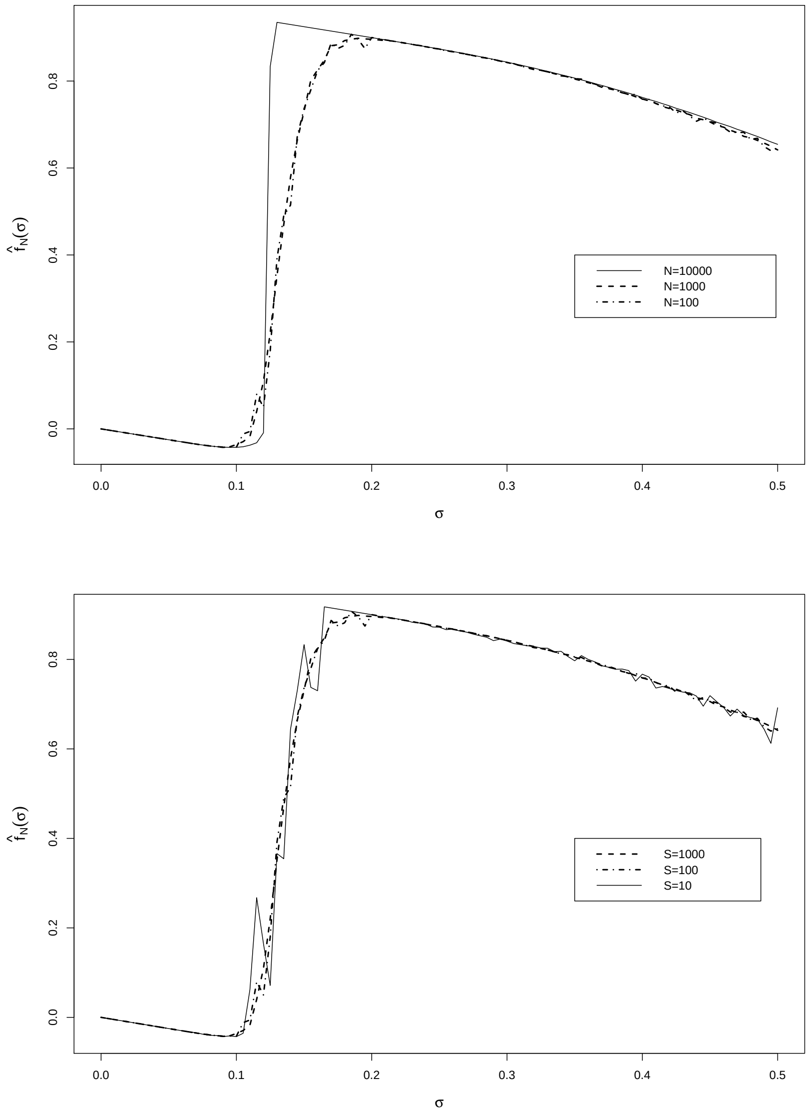

Figure 4: The objective function $f_N(\sigma)$. Top panel: for fixed $S = 100$, the cases $N = 100, 1000, 10000$ are shown. Bottom panel: for fixed $N = 1000$, the cases $S = 10, 100, 1000$ are displayed.

The bottom panel illustrates the role of S, for given $N = 1000$: when $S = 10$, admittedly a case in which the true expectation is poorly estimated, the function $\hat{f}_N(\sigma)$ is noisy and displays jumps that would make any maximization effort hard or vain. On the contrary, when $S = 100$ or $S = 1000$, the graphs are smoother and almost overlapped; this seems to suggest that such numbers of sample observations are appropriate to reasonably approximate σ_N^*. Practically speaking, some balance must be stricken and, if N is to be increased to explore asymptotics, S should be limited to avoid an excessive computational burden.

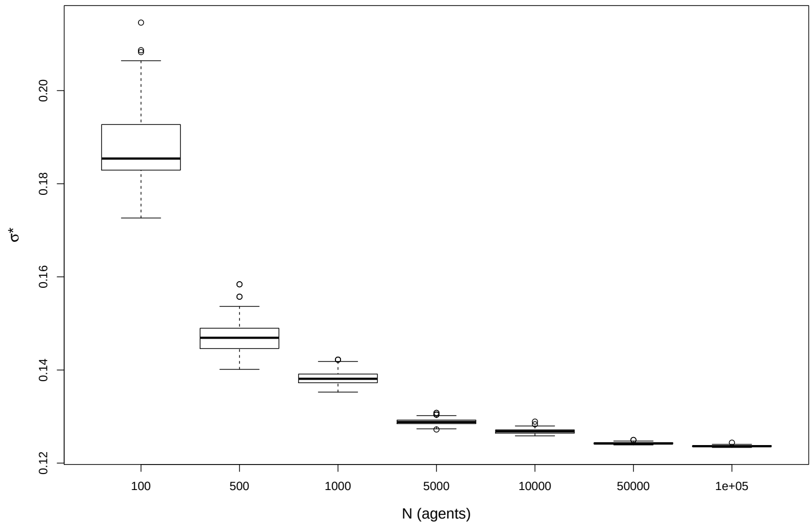

Figure 5: Boxplots of M maximizers ${}_m\hat{\sigma}_N^*, m = 1, ..., M$, for $N = 100, 500, 1000, 5000, 10000, 50000, 100000$ and fixed $S = 100$.

Based on the above argument, we set $S = 100$ and solve $M = 100$ independent instances of the optimization problems, for increasing values of $N = 100, 500, 1000, 5000, 10000, 50000, 100000$. Figure 5 shows the boxplots of ${}_m\sigma_N^*, m = 1, ..., M$, in the various cases. The boxes depict interquartile ranges (IQR); the whiskers extend up to 1.5 IQR to provide evidence of outliers (circles), if any; the horizontal line is the median value. While, for small values of N, the true maximizer $\sigma_c \approx 0.122$, referred to in Theorem 3.3, is overestimated and there are some outlying results, as N increases most of the maximizers of the (stochastic) problems are quite close to the correct result and, clearly, boxplots support the convergence rigorously proven in the previous section. In other words, if N is too small, the maximizers obtained by numerical optimization inaccurately span a wide interval, that can be roughly located at $[0.17, \ 0.21]$. As N is increased, the numerical results are increasingly close to σ_c and are no longer dispersed. The plotted data were obtained using the routine `optimize` in `R`, R Core Team [2018], and computations took about 18 minutes on a 51-cores 2.2 GHz Linux cluster set up with the library `parallel`. It may be worth noticing that, say, the boxplot for

$N = 100000$ required to sample an order of $M \cdot S \cdot N = 100 \cdot 100 \cdot 100000 = 10^9$ individual thresholds.[15]

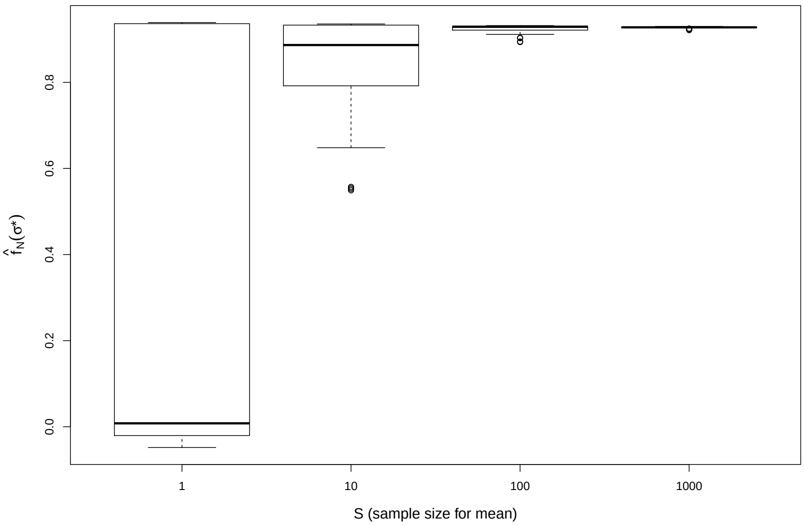

Figure 6: Boxplots of M maxima ${}_m\hat{f}^*_N, m = 1, ..., M$, for $S = 1, 10, 100, 1000$ and fixed $N = 1000$.

Not surprisingly, the value of S also has a conspicuous impact on the accuracy of the M solutions of the problems that are numerically found in our simulations. Figure 6 depicts, for fixed $N = 1000$, the boxplots of the maxima ${}_m\hat{f}^*_N, m = 1, ..., M$, for different values of S: for $S = 1$ the computed results are essentially unreliable and are spread over a large interval. Notice that the boxplot relative to $S = 1$ well represents the ill-posedness depicted in Figure 3: unless the mean operator is used, moving down along the right side of the diagram by increasing N does not remove the discontinuity of $f(\sigma)$. As a consequence, numerically computed maxima are most often very far from $\tilde{r}_c$. On the opposite, by increasing S, i.e., by taking the mean on the left side of the diagram, and then maximizing the objective function f_N such that $f_N \xrightarrow{\Gamma} f_\infty$, produces a sequence of pairs of real numbers $(\sigma^*_N, \rho_N(\sigma^*_N))$ which converges to $(\sigma_c, \tilde{r}_c)$.

Improving the accuracy of the sample mean by using larger S reduces the errors and visually confirms again that S in the range 100–1000 appears to be numerically satisfactory. Somewhat more deeply, the figure portrays again the importance of the smoothing effect provided by the average operator:

[15]As `optimize` typically evaluates the objective function about 20 times, the sampled thresholds exceed 10^{10}.

on the first hand, problems can be more accurately solved; but, on the other hand, the ill-posedness of the original formulation is removed and tractability is obtained in the form of Γ-convergence.

5 Conclusions

We have studied a class of "optimization problems over a dynamical system" and have shown that care is needed to avoid ill-posedness. This was exemplified using the well-known riot example discussed in Granovetter [1978]: intuitively, at some critical value for one parameter, the resulting endogenous equilibrium can abruptly jump due to small and natural sampling variability in the activation thresholds of the agents. More formally, in the presence of a saddle-node bifurcation of the large-scale deterministic dynamics, some equilibria disappear and the objective of the policy-maker turns out to be discontinuous, breaking down conventional optimization approaches. The problem cannot be removed by merely increasing, or taking the limit of, the number N of agents and solving the deterministic version of the problem, due to the tangency condition depicted in Figure 1.

Hence, we considered a sequence of finite-dimensional problems, where the steady states of the dynamical system are random variables and we proved that the objective functions of such problems are now continuous thanks to the *regularizing effect* of the expectation operator. Moreover, the sequence of the maximizers converges to the critical value of the parameter, i.e. the value where the bifurcation happens, and the sequence of maxima converges to the supremum of the (ill-posed) original problem. Technically speaking, this convergence applies in a Γ-sense and we believe that ours may be among the very few practical applications of this methodology to economically-relevant problems. Specifically, the sequence of the functions f_N that are maximized Γ-converges towards f_∞ which is proven to be the upper-semicontinuous envelope of the original objective function f. Interestingly, f_∞ and f only differ at the value σ_c, which is the limit of the sequence of maximizers, and, in turn, the point where the social planner (and we, as rational and conscious external observers) expected a maximum to exist in the ill-posed original problem. In this respect, our approach endows the social planner with a mathematically precise argument to identify the unique optimal policy.

From the theoretical point of view, the use of the expectation operator appears to be crucial in the definition of a sequence of problems that converges to a regularized right-continuous objective function, where the maximum is now well-defined and coincides with the supremum of the original problem.

Our numerical results make also clear that the use of an expectation is not simply an astute device to get some usable approximated result but, more fundamentally, poses the basis for the Γ-convergence (in the number of agents N) to the theoretical model. This, together with a proper *Law of Large Numbers* in S, the sample size of the estimator of the expectation, ensures convergence also of the numerical simulations.

We have demonstrated, in this specific case, that sample sizes of the order of one hundred to one thousand provide enough smoothing to get reasonably accurate numerical results (with the help of a multi-core processor and parallel computations). We feel that our exemplification shows the need to incorporate proper sampling schemes in similar problems or whenever, say, a numerical treatment is the only viable option due to the lack of closed-form expressions for the steady state of the dynamical systems.

A Proofs

In this appendix, we collect all proofs.

Proof of Lemma 2.1

We first state a technical classical result on bifurcation theory of dynamical system.

Lemma A.1 (saddle-node bifurcation). *Consider the dynamical system*

$$x(t+1) = F(x(t);\sigma); \quad x(0) = x_0 \in [0,1] \tag{14}$$

where $x \in [0,1]$ and $F(\cdot\,;\sigma)$ is a continuous distribution function with standard deviation $\sigma \geq 0$, admitting a unimodal density function. The set of steady states of (14) *is given by the solutions to*

$$x = F(x;\sigma). \tag{15}$$

There exists a threshold level σ_c, such that:

i) if $\sigma > \sigma_c$, (15) *admits a unique solution x;*

ii) if $\sigma = \sigma_c$, (15) *admits two solutions $x^l < x^h$;*

iii) if $\sigma < \sigma_c$, (15) *admits three solutions $x^l < x^m < x^h$.*

If the solution x is unique, it is a globally stable attractor for the dynamical system in (14)*. In case of multiple equilibria, if $x_0 < x^m$, then*

$$\lim_{t\to\infty} x(t) = x^l.$$

On the opposite, if $x_0 > x^m$*, then*

$$\lim_{t\to\infty} x(t) = x^h.$$

Back to the proof of Lemma (2.1), we are exactly in this situation, since F is Gaussian. The graph of the distribution function F intersects the bisector either three times, twice or once. A visual representation of the three different cases is reported in Figure 1. When $\sigma = \sigma_c$, we are in the tangency situation: the graph of F intersects the bisector line at one point x^l, when the curves are tangent, and in one second point x^h, when they are not. If the expected value μ of the distribution is such that $\mu < 1/2$, then it is easy to see that $x^l < 1/2 < x^h$. Consider finally that we take $r_0 = 0$ so that the dynamical system always converges to x^l, i.e., the smallest among the solutions to (15). Let us call this equilibrium $\rho(\sigma)$. The previous lemma shows that the map $\sigma \mapsto \rho(\sigma)$ is continuous on $[0, \sigma_c)$ and on $(\sigma_c, +\infty)$ with an upward jump at σ_c. This proves left-continuity. It remains to show that $f(\sigma) = \rho(\sigma) - k\sigma$ is bounded from above and that the supremum of f is exactly $\tilde{f}$ defined as $\lim_{\sigma\to\sigma_c^+} f(\sigma)$. It is convenient to separate the study on the two intervals $[0, \sigma_c]$ and $(\sigma_c, +\infty)$. Concerning the latter, we show that, on this interval, $\sigma \mapsto \rho(\sigma)$ is decreasing. As remarked above, when $\sigma > \sigma_c$ there exists a unique x solving (15), with $x > \mu$. Moreover, still for $\sigma > \sigma_c$ and $x > \mu$, $F(x;\sigma)$ is concave and increasing. Let $\sigma_1, \sigma_2 \in (\sigma_c, +\infty)$ be such that $\sigma_1 < \sigma_2$. Then $F(\cdot;\sigma_1)$ and $F(\cdot;\sigma_2)$ satisfy

$$F(x;\sigma_2) < F(x;\sigma_1) \quad \forall\, x > \mu. \tag{16}$$

Let us call ξ_1 the unique solution to $F(x;\sigma_1) = x$. Then, by (16),

$$F(\xi_1;\sigma_2) < F(\xi_1;\sigma_1) = \xi_1.$$

Since $F(\cdot;\sigma_2)$ is increasing, there exists $\xi_2 < \xi_1 : F(\xi_2;\sigma_2) = \xi_2$. Note that $\xi_1 = \rho(\sigma_1)$ and $\xi_2 = \rho(\sigma_2)$. Therefore, $\rho(\sigma)$ is decreasing on $(\sigma_c, +\infty)$ as so also f is decreasing on that interval. Concerning the interval $[0, \sigma_c]$, a similar concavity argument shows that, in this case, $\sigma \mapsto \rho(\sigma)$ is increasing up to a level $\rho(\sigma_c)$ such that $\lim_{\sigma\to\sigma_c^+} \rho(\sigma) := \tilde{r} > \rho(\sigma_c)$. Since, by assumption,

$$k < k^{th} := \frac{\tilde{r} - \rho(\sigma_c)}{\sigma_c},$$

we have that, on $[0, \sigma_c]$,

$$k < \frac{\tilde{r} - \rho(\sigma_c)}{\sigma_c} < \frac{\tilde{r} - \rho(\sigma_c)}{\sigma_c - \sigma} < \frac{\tilde{r} - \rho(\sigma)}{\sigma_c - \sigma},$$

where the latter is due to the monotonicity of ρ. Thus,

$$k < \frac{\tilde{r} - \rho(\sigma)}{\sigma_c - \sigma} \iff k(\sigma_c - \sigma) < \tilde{r} - \rho(\sigma) \iff \rho(\sigma) - k\sigma < \tilde{r} - k\sigma_c := \tilde{f}.$$

Summarizing, we have proved that: (i) f is left-continuous with a discontinuity at σ_c; (ii) the function f is bounded from above and admits a finite supremum $\tilde{f}$ which is not a maximum. □

Proof of Theorem 3.3

We first prove a technical lemma.

Lemma A.2. *For each N, consider R_N as defined in* (11)*. R_N is a measurable and bounded function of the finite sample $\tilde{X} \equiv (X_1, \dots, X_N)$. Moreover, the function $\sigma \mapsto \rho_N(\sigma)$ is continuous.*

Proof.

$$\rho_N(\sigma) = \mathbb{E}^{\sigma}[R_N] = \int_{supp(\tilde{X})} R_N(\tilde{x})\, d\tilde{F}(\tilde{x}, \sigma), \tag{17}$$

where $\tilde{F}(\tilde{x};\sigma) = \Pi_{i=1}^{N} F(x_i, \sigma)$ and where $\tilde{x} = (x_1, \dots, x_N)$ is a realization of the sample $\tilde{X}$. Note that the integrand function $R_N(\tilde{x})$ is a measurable and bounded function of the sample. As a consequence, the integral is well defined; moreover, it is continuous in σ as soon as F is continuous in σ. □

Back to the statement of Theorem 3.3, to prove the well-posedness of the problem (P2), simply note that, thanks to Lemma A.2, $f_N(\sigma) = \rho_N(\sigma) - k\sigma$ is continuous, $\rho_N(\sigma)$ is bounded and, finally, $\lim_{\sigma\to\infty} f_N(\sigma) = -\infty$. Therefore, the objective function is continuous and bounded from above, hence it admits a maximum point and a maximizer σ_N^*.

For the second part of the statement we use the tool of the Γ-convergence which, under suitable conditions, implies convergence of minimum values and minimizers. Actually, so far we have considered the maximization problem (P2). To align our reasoning with the standard literature in analysis, from now on, we take the equivalent minimization problem defined as

$$\min_{\sigma} \overline{f}_N(\sigma) \tag{18}$$

where

$$\overline{f}_N(\sigma) = -\rho_N(\sigma) + k\sigma = -f_N(\sigma)$$

is the sequence of the objective functions of the problem (18). It is easy to see that the sequence of objective functions $(\overline{f}_N)_N$ defined on $\mathbb{R}^+$ is *equi-mildly* *coercive.*[16] Moreover, given the function $\overline{f}(\sigma) = -\rho(\sigma) + k\sigma$, we consider its lower-semicontinuous envelope $sc\overline{f}$ (see Definition B.4), that is, for every $\sigma \in \mathbb{R}^+$

$$sc\overline{f}(\sigma) = \begin{cases} -\rho(\sigma) + k\sigma & \text{if } \sigma \neq \sigma_c \\ -\tilde{r}_c + k\sigma_c & \text{if } \sigma = \sigma_c \end{cases} \tag{19}$$

with $-\tilde{r}_c := -\lim_{\sigma \to \sigma_c^+} \rho(\sigma)$. As showed in Proposition B.5 $sc\overline{f}(\sigma)$ is the $\Gamma-$limit of $\overline{f}_N(\sigma)$, i.e., renaming $sc\overline{f}(\sigma) = \overline{f}_\infty$ one has that

$$\overline{f}_N(\sigma) \xrightarrow{\Gamma} \overline{f}_\infty.$$

Then, using Theorem B.8 in our framework, we get

$$\exists \min_{\mathbb{R}^+} \overline{f}_\infty(\sigma) = \lim_{N \to +\infty} \inf_{\mathbb{R}^+} \overline{f}_N(\sigma). \tag{20}$$

Moreover, since all functions $\overline{f}_N$ admit a minimizer $\overline{\sigma}_N^*$ (which exists in virtue of Lemma A.2) then, up to subsequences, $\overline{\sigma}_N^*$ converges to a minimum point of $\overline{f}_\infty$. According to (19), the only minimizer of $\overline{f}_\infty$ is σ_c. Hence, by (20), it follows that

$$\lim_{N \to +\infty} -\rho_N(\overline{\sigma}_N^*) + k\overline{\sigma}_N^* = -\hat{r}_c + k\sigma_c. \tag{21}$$

Accordingly,

$$\overline{\sigma}_N^* \to \sigma_c. \tag{22}$$

Therefore, from (21) and (22) and back to the maximization problem (P2), it follows that the sequence of objective functions f_N of the auxiliary problem (P2) Γ-converges to f_∞, the upper-semicontinuous envelope of the objective function of the original problem (P1). As a consequence, $\rho_N(\sigma_N^*)$ converges to $\tilde{r}_c$ and σ_N^* converges to σ_c. □

B Some basics of Γ-convergence

In this section, we introduce some abstract notions and results on Γ-convergence. We start recalling the concepts of lower and upper limits and of lower-semicontinuous functions to introduce the definition of Γ-limit. We also define the lower-semicontinuous envelope of a function, its link with Γ-limit and we provide an example of computation of Γ-limit by noting how this can be different from the pointwise limit. Finally, we show that under suitable

[16]See Definition B.7. In particular, our sequence $(\overline{f}_N)_N$ is equi-mildly coercive since all functions $\overline{f}_N$ are bounded from below. Taking $\tilde{\sigma}$ large enough, all infima are reached in $[0, \tilde{\sigma}]$ which is compact.

conditions Γ-convergence implies convergence of minimum values and minimizers.

From now on, unless otherwise specified X will be a metric space equipped with the metric d.

Definition B.1. *Let $f : X \to \overline{\mathbb{R}}$. We define the lower limit (liminf for short) of f at x as*

$$\begin{aligned}\liminf_{y\to x} f(y) &= \inf\{\liminf_j f(x_j) : x_j \in X, x_j \to x\}\\ &= \inf\{\lim_j f(x_j) : x_j \in X, x_j \to x, \exists \lim_j f(x_j)\},\end{aligned}$$

and the upper limit (limsup for short) of f at x as

$$\begin{aligned}\limsup_{y\to x} f(y) &= \sup\{\limsup_j f(x_j) : x_j \in X, x_j \to x\}\\ &= \sup\{\lim_j f(x_j) : x_j \in X, x_j \to x, \exists \lim_j f(x_j)\}.\end{aligned}$$

Definition B.2. *A function $f : X \to \overline{\mathbb{R}}$ is lower semicontinuous at $x \in X$, if for every sequence (x_j) converging to x we have*

$$f(x) \leq \liminf_j f(x_j), \tag{23}$$

or, in other words, $f(x) = \min\{\liminf_j f(x_j) : x_j \to x\}$. We will say that f is lower semicontinuous on X if it is l.s.c. at all $x \in X$.

Definition B.3 (Γ-convergence)**.** *A sequence $f_j : X \to \overline{\mathbb{R}}$ Γ-converges in X to $f_\infty : X \to \overline{\mathbb{R}}$ if for all $x \in X$ we have*

(i) (liminf inequality) for every sequence (x_j) converging to x

$$f_\infty(x) \leq \liminf_j f_j(x_j); \tag{24}$$

(ii) (limsup inequality) there exists a sequence (x_j) converging to x such that

$$f_\infty(x) \geq \limsup_j f_j(x_j). \tag{25}$$

The function f_∞ is called the Γ-limit of (f_j), and we write $f_\infty = \Gamma\text{-}\lim_j f_j$.

Definition B.4. *Let $f : X \to \overline{\mathbb{R}}$ be a function. Its lower semi-continuous envelope scf is the greatest lower-semicontinuos function not greater than f, that is, for every $x \in X$*

$$scf(x) = \sup\{g(x) : g \text{ l.s.c.}, g \leq f\}.$$

Proposition B.5. *It shows that* $scf(x) = f_\infty(x)$*.*

Proof. See Proposition 1.31 of Braides [2002].

Here below, we report an example which highlights the different roles of the limsup and liminf inequalities. It is also useful to visualize in a simple case of a sequence of real functions the difference between the classical pointwise (or uniform) limit and the Γ-limit.

Example B.6. *Let* $f_j(t)$ *be a sequence of function, where*

$$f_j(t) = \begin{cases} \pm 1 & \textit{if} \quad t = \pm 1/j, \\ 0 & \quad \textit{otherwise.} \end{cases}$$

Note that $f_j \to 0$ *pointwise, but* $\Gamma\text{-}\lim_j f_j = f_\infty$*, where*

$$f_\infty(t) = \begin{cases} -1 & \textit{if} \quad t = 0, \\ 0 & \textit{if} \quad t \neq 0. \end{cases}$$

Indeed, the sequence f_j *converges pointwise (and hence also* Γ*-converges) to* 0 *in* $\mathbb{R} \setminus \{0\}$*, while the optimal sequence for* $t = 0$ *is* $t_j = -1/j$*, for which* $f_j(t_j) = -1$*.*

Definition B.7 (Coerciveness conditions)**.** *A function* $f : X \to \overline{\mathbb{R}}$ *is mildly coercive if there exists a non-empty compact set* $K \subset X$ *such that* $\inf_X f = \inf_K f$*. A sequence* (f_j) *is equi-mildly coercive if there exists a nonempty compact set* $K \subset X$ *such that* $\inf_X f_j = \inf_K f_j$ *for all* j*.*

Theorem B.8. *Let* (X, d) *a metric space, let* (f_j) *be a sequence of equi-mildly coercive functions on* X*, and let* $f_\infty = \Gamma\text{-}\lim_j f_j$*; then*

$$\exists \min_X f_\infty(x) = \lim_{j \to +\infty} \inf_X f_j(x). \tag{26}$$

Moreover, if all functions f_j *admit a minimizer* x_j^* *then, up to subsequences,* x_j^* *converges to a minimum point of* f_∞*.*

Proof. See Theorem 1.21 and Remark 1.22 in Braides [2002].

Acknowledgements

We thank Marco LiCalzi for insightful discussions and comments. Paolo Dai Pra raised our attention on Γ-convergence. The work was partially funded by the ITN ExSIDE European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agreement No 721846.

References

P.M. Allen, M. Sanglier. A dynamic model of growth in a central place system, *Geographical Analysis*, 11(3), 256-272, 1979.

E. Barucci, M. Tolotti. Social interaction and conformism in a random utility model. *Journal of Economic Dynamics and Control*, 36(12), 1855-1866, 2012.

E. Barucci, M. Tolotti. Identity, reputation and social interaction with an application to sequential voting. *Journal of Economic Interaction and Coordination*, 7(1), 79-98, 2012.

F.M. Bass. A new product growth for model consumer durables. *Management science*, 15(5), 215-227, 1969.

L. Blume, S. Durlauf. Equilibrium concepts for social interaction models. *International Game Theory Review*, 5(03), 193-209, 2003.

A. Braides. Γ-convergence for beginners, *Oxford Lecture Series in Mathematics and its Applications*, 22, Oxford University Press, Oxford, 2002.

W.A. Brock, S. Durlauf. Discrete choice with social interactions. *The Review of Economic Studies*, 68(2), 235-260, 2001.

G. Carlier. A general existence result for the principal-agent problem with adverse selection, *J. Math. Econom.*, 35(1), 129-150, 2001.

S.N. Ethier, T.G. Kurtz. Markov processes: characterization and convergence (Vol. 282). John Wiley & Sons, 2009.

M. Ghisi, M. Gobbino. The monopolist's problem: existence, relaxation, and approximation, *Calc. Var. Partial Differential Equations*, 24(1), 111-129, 2005.

F. Ghoulmie, R. Cont, and J-P. Nadal. Heterogeneity and feedback in an agent-based market model, *J. Phys. Condens. Matter*, 17, S1259–S1268, 2005.

M. Granovetter. Threshold models of collective behavior. American journal of sociology, 83(6), 1420-1443, 1978.

P. Milgrom, J. Roberts. Comparing equilibria. The American Economic Review, 441-459, 1994.

P. Monteiro, F. H. Page Jr. Optimal selling mechanisms for multiproduct monopolists: incentive compatibility in the presence of budget constraints, *J. Math. Econom.* 30(4), 473-502, 1998.

J-P. Nadal, D. Phan, M. B. Gordon, J. Vannimenus. Multiple equilibria in a monopoly market with heterogeneous agents and externalities, *Quant. Finance*, 5(6), 557–568, 2005.

M. Gordon, J-P. Nadal, D. Phan, V. Semeshenko. Entanglement between demand and supply in markets with bandwagon goods, *J. Stat. Phys.*, 151, no. 3-4, 494–522, 2013.

R. Peres, E. Muller, V. Mahajan. Innovation diffusion and new product growth models: A critical review and research directions. *International journal of research in marketing*, 27(2), 91-106, 2010.

R Core Team. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. `https://www.R-project.org/`, 2018.

J.C. Rochet, P. Choné. Ironing, sweeping and multidimensional screening, *Econometrica*, 66, 783-826, 1998.

S. H. Strogatz. Nonlinear dynamics and chaos. With applications to physics, biology, chemistry, and engineering. Second edition. Westview Press, Boulder, CO, 2015.

T. C. Schelling. Dynamic Models of Segregation, *Journal of Mathematical Sociology*, 1(2), 143-186, 1971.