

Spatial extreme values: variational techniques and stochastic integrals

Nicolás García Trillos* Ryan Murray† Daniel Sanz-Alonso‡

August 2, 2018

Abstract

This work employs variational techniques to revisit and expand the construction and analysis of extreme value processes. These techniques permit a novel study of spatial statistics of the location of minimizing events. We develop integral formulas for computing statistics of spatially-biased extremal events, and show that they are analogous to stochastic integrals in the setting of standard stochastic processes. We also establish an asymptotic result in the spirit of the Fisher-Tippett-Gnedenko theory for a broader class of extremal events and discuss some applications of our results.

1 Introduction

Extreme value theory is the branch of probability theory and statistics concerned with the study of extreme deviations from the median behavior. A fundamental problem is to characterize the asymptotic distribution of

$$M_n = a_n \min_{i=1\dots n} (X_i - b_n),$$

where the X_i are i.i.d. random variables. A complete account of the possible limit distributions of this type of process is given by the Fisher-Tippett-Gnedenko theory (see e.g. [Res13]), which also explains the role of the normalizing sequences $\{a_n\}$ and $\{b_n\}$. For instance if the X_i are strictly positive with unit density at zero and $a_n = n, b_n = 0$, then the M_n converge in law to the exponential distribution with unit rate. This classic theory and some related stochastic processes will be reviewed in Subection 1.2. In many modern settings it is desirable to not only understand the distribution of minima, but also of the index (or “spatial location”) of minimizers; this paper investigates some questions arising in such settings. Suppose, as an illustration, that a user sends a request to a group of n spatially-distributed servers with distinct locations $x_1, \dots, x_n \in D \subset \mathbb{R}^d$. The user and the parties running the servers will certainly be interested in waiting time for the first response to occur, but also will likely be interested in identifying the location of the first responding server. In order to model such a system, we will suppose that the response time of the servers is given by

$$f_n(x) := \begin{cases} \frac{n}{\lambda(x_i)} \xi_i + g(x_i), & \text{for } x = x_i, \quad i = 1, \dots, n, \\ +\infty, & \text{otherwise,} \end{cases} \quad (1.1)$$

*Department of Statistics, University of Wisconsin.

†Department of Mathematics, Pennsylvania State University

‡Department of Statistics, University of Chicago.

where $g(x_i)$ represents a deterministic communication latency (i.e. the amount of time it takes the message to reach the i -th server) that is *not* necessarily spatially uniform, the ξ_i are i.i.d. exponential variables, and the $\lambda(x_i)$ represents the processing speed of the server at location x_i (which could be related to either hardware or workload). We have chosen the scaling n for the random exponential term so that in the large n limit it has comparable size to the deterministic latency. In a way that will be made precise, this scaling ensures that, for large n , the overall processing rate of all servers is of order 1 and given by the function λ . The latency imposes a relatively weak deterministic spatial correlation structure and models a scenario where the processing power needed for the request is high and the communication overhead is low. The locations x_i of the servers will be assumed to be distributed according to a probability density ρ in a bounded domain $D \subset \mathbb{R}^d$ so that, in general, the locations will not be uniformly distributed. In this framework, it is natural to study the distribution of the first response time, along with the distribution of the *location* of the first server making a response.

A first contribution of this paper is to provide integral representation formulas, which to our knowledge are novel, for the distribution of mins and argmins. These formulas may be of use in various inference problems, some of which will be outlined in Section 1.3. We will be mostly interested in the limit problem as $n \rightarrow \infty$ and a second contribution is to introduce an appropriate analytical setting in which to study such a limit using modern tools from the calculus of variations [DM93]. The response time functions f_n will be viewed as random variables with values in the space of lower semi-continuous functions $\mathcal{S}(D)$ (a precise definition of this space, and the associated topology of Γ -convergence, will be given in Section 2.1). Intuitively, $\mathcal{S}(D)$ is a natural choice of function space, as it permits the evaluation of minima and minimizers both in the discrete case with finite n and in the limiting case, $n \rightarrow \infty$. Moreover, the functions f_n above are immediately lower semi-continuous, as they are defined to be $+\infty$ except at the x_i . Finally, a third contribution of this paper is to draw several analogies between extreme value processes and standard stochastic processes. We show that, if $g \equiv 0$, the large n distribution of mins and argmins of the response times f_n is governed by the distribution of mins and argmins of a limiting extreme value process W_λ that can be seen as an analog of Brownian motion. The process W_λ has been studied before, but here our focus is on the spatial statistics of minimizers, which differs from much of the previous literature –see Subsection 1.2. In the case of general g , evaluating local minima is analogous to computing stochastic integrals of g , with extreme value processes playing the role of Brownian motion.

The remainder of the introduction goes as follows. In Subsection 1.1 we state our main results. We review related work in Subsection 1.2, and applications and future lines of research in Subsection 1.3. We close with an outline of the paper in 1.4.

1.1 Main results

The main results of this paper are concerned with the random process W_λ , which will serve as a limit of the process described above in the case $g \equiv 0$. Here and throughout D denotes a bounded domain in \mathbb{R}^d , \bar{D} its closure, and $\mathcal{S}(\bar{D})$ the space of lower semi-continuous function on \bar{D} (see Section 2 for more details on this space).

Definition 1.1. *Let $\lambda: \bar{D} \rightarrow (0, \infty)$ be a continuous function. We define W_λ to be the random object with values in $\mathcal{S}(\bar{D})$ satisfying the following properties:*

1. *For every closed set C the random variable $\min_{x \in C \cap \bar{D}} W_\lambda(x)$ is exponentially dis-*

tributed with rate

$$\lambda_C := \int_{C \cap \bar{D}} \lambda(x) dx.$$

2. For any finite collection of disjoint and closed sets C_1, \dots, C_k , the random variables

$$\min_{x \in C_1} W_\lambda(x), \dots, \min_{x \in C_k} W_\lambda(x)$$

are independent.

We will show in Propositions 2.9 and 2.10 that in fact the random process W_λ exists and is well defined. Furthermore, we will show in Proposition 2.19 that after adding a deterministic function g to W_λ , the resulting random function has well-defined, unique k -argmins (as defined in Definition 2.14). Our first main result establishes a representation formula for the joint distribution of the first k -argmins.

Theorem 1.2. *Let $g \in \mathcal{S}(\bar{D})$ (see Section 2.1) and let W_λ be as in Definition 1.1. Then, (with probability 1) the random function $W_\lambda + g$ has well-defined, unique k -argmins for every $k \in \mathbb{N}$ (see Definition 2.14). In addition, the joint distribution of the first k -argmins $(X^{(1)}, \dots, X^{(k)})$ is given by the density function:*

$$\begin{aligned} \frac{\mathbb{P}(X^{(1)} \in dx_1, \dots, X^{(k)} \in dx_k)}{dx_1 \dots dx_k} &= \int_{\mathbb{R}^{k-1}} \Phi(x_k, (g - r_{k-1})^+) \Psi(x_1, r_1, g) \\ &\quad \times \prod_{j=2}^{k-1} \Psi(x_j, r_j - r_{j-1}, (g - r_{j-1})^+) dr_1 \dots dr_{k-1}, \end{aligned} \quad (1.2)$$

with

$$\begin{aligned} \Phi(z, h) &:= \int_{h(z)}^{\infty} \exp\left(-\bar{\lambda} \int_{-\infty}^t H_{\lambda, h}(s) ds\right) dt, \\ \Psi(z, t, h) &:= \chi_{t > h(z)} \exp\left(-\bar{\lambda} \int_{-\infty}^t H_{\lambda, h}(s) ds\right), \\ H_{\lambda, h}(s) &:= \mu_\lambda(\{x \in D : h(x) \leq s\}), \\ \bar{\lambda} &:= \int_D \lambda(x) dx, \\ \mu_\lambda(E) &:= \frac{1}{\bar{\lambda}} \int_E \lambda(x) dx. \end{aligned} \quad (1.3)$$

Our second main result establishes that $W_\lambda + g$ serves as an appropriate limiting process for a wide class of discrete processes, including the f_n above.

Theorem 1.3. *Suppose that x_1, \dots, x_n, \dots is a sequence of points in D for which the empirical measures*

$$\frac{1}{n} \sum_{i=1}^n \delta_{x_i}$$

converge weakly, as $n \rightarrow \infty$, towards the measure $\rho(x) dx$, where $\rho: D \rightarrow \mathbb{R}$ is a continuous density function. Let $\lambda: D \rightarrow \mathbb{R}$ be a continuous function and let $g: \bar{D} \rightarrow \mathbb{R}$. Let $\xi_1, \dots, \xi_n, \dots$ be independent, non-negative random variables (not necessarily identically distributed) whose cumulative distribution functions F_i satisfy

$$F_i(t) = t + O(t^2) \quad (1.4)$$

near zero (uniformly in i). Define the random function f_n by

$$f_n(x) = \begin{cases} \frac{n}{\lambda(x_i)} \xi_i + g(x_i), & \text{if } x = x_i, \quad i = 1, \dots, n, \\ +\infty, & \text{otherwise.} \end{cases}$$

Then, the joint distribution of the first k argmins of f_n converges, as $n \rightarrow \infty$, towards the joint distribution of the first k argmins of $W_{\tilde{\lambda}} + g$, where

$$\tilde{\lambda}(x) = \lambda(x)\rho(x), \quad x \in D.$$

The proofs of Theorems 1.2 and 1.3 will be given in Section 3. An immediate corollary of these theorems is the following:

Corollary 1.4. *In the setting of Theorem 1.3, the joint distribution of the first k argmins of f_n converges, as $n \rightarrow \infty$, towards the distribution given by (1.2) with λ replaced by $\lambda(x) = \lambda(x)\rho(x)$.*

Remark 1.5. *In most proofs below the random variables ξ_i are assumed to be exponentially distributed with unit rate and the points x_i to be uniformly distributed. This assumption can be made without loss of generality, as the rates of the ξ_i and fluctuations in ρ may be absorbed into λ . Moreover, the assumption that the ξ_i are exponential may be relaxed in the limit as $n \rightarrow \infty$ (see e.g. Theorem 1.3). Second, the Fisher-Tippett-Gnedenko theory provides a broader class of possible limiting distributions (analogous to α -stable distributions in the classical central limit theory), which would not require such a specific behavior of the F_i near zero. We believe that it is of interest to extend our results to that broader family, but such an extension is beyond the scope of this work.*

1.2 Related work

Extreme value theory and extremal processes is a classical branch of probability and statistics. While we do not attempt to give an exhaustive account of the field, we will highlight some issues that are relevant to our work. We will also draw parallels with other limit theorems and stochastic processes.

One of the first questions in probability theory is to understand the limits of combinations of independent random variables. In the context of the central limit theorem, one considers

$$\lim_{n \rightarrow \infty} n^{-1/2} \left(\sum_{i=1}^n X_i - n\mathbb{E}(X_i) \right).$$

The central limit theorem states that, for i.i.d. X_i with finite variance, the limit is a normally-distributed random variable. In the context of extreme values, given a sequence $\{X_i\}$ of i.i.d. variables one similarly considers

$$\lim_{n \rightarrow \infty} a_n \left(\min_{i=1 \dots n} X_i - b_n \right).$$

Here a_n and b_n are normalizing constants, analogous to $n^{-1/2}$ and $n\mathbb{E}(X_i)$ in the central limit theorem case. The Fisher-Tippett-Gnedenko Theorem [Gne43] completely specifies possible limits for this process. A detailed description of this theory can be found in [LLR83] or [Res13]. In this paper we focus on the case where X_i is positive, and has a density function which is positive at zero.

Subsequently, various authors [Dwa64] [Lam64] studied the distribution of the “ k -records”

$$M_n^k = a_n \left(\underset{i=1\dots n}{k\text{-min}} X_i - b_n \right),$$

where $\underset{i=1\dots n}{k\text{-min}}$ denotes the value of the k -th smallest element in the collection. Studying k -th mins is an important branch of extreme value theory that we pursue and extend here by characterizing the spatial distribution of the k -th minima.

Returning to the discussion of first mins, a natural object is the rescaled process

$$M_n(t) = a_n \left(\min_{i < nt} X_i - b_n \right).$$

that tracks mins over time. Taking the limit $\lim_{n \rightarrow \infty} M_n(t) =: M(t)$ yields so-called *extremal processes* [Res13]. The construction of such processes is completely analogous to the construction of Brownian motion by Lévy, where sums are replaced by mins.

In standard stochastic processes, a central concept is the family of *stable stochastic processes*, which are invariant under linear combinations. In the context of extreme values one instead considers the family of *max-stable processes*, which are stochastic processes that satisfy

$$\max_{i=1\dots r} M^{(i)}(t) \stackrel{d}{=} rM(t),$$

where the $M^{(i)}$ are independent copies of M and $\stackrel{d}{=}$ denotes equality in distribution. A significant literature studies these processes by means of Poisson processes on the plane [Pic71], and gives a spectral representation of these processes [dH84]. Min-stable processes (namely processes M_t for which $\frac{1}{M_t}$ is max-stable) were studied in [dHP84]. Again, much of this literature focuses on the distribution of minima and not on the locations at which minima occur.

The previous objects can be easily generalized by evaluating the minimum value over some Borel set $A \subset \mathbb{R}$,

$$M_n(A) = a_n \left(\min_{\frac{\cdot}{n} \in A} X_i - b_n \right).$$

Note that letting $A = [0, t)$ recovers the previously-defined process $M_n(t)$. Again, one may then consider the limit

$$M(A) := \lim_{n \rightarrow \infty} M_n(A).$$

The set functions M_n, M are called *inf-measures*. In the context of maximization the analog is known as a *sup-measure*.

In this paper we work at the level of lower semi-continuous functions, that is, we study directly the functions $\lim_{n \rightarrow \infty} nX_{\frac{\cdot}{n}}$ as opposed to the minimum value that they take. This seems more natural here, as it allows us to study simultaneously both minimizers and minima. The approach via lower semi-continuous functions can be shown to be equivalent to that of inf-measures. Namely, given a inf-measure M , we can define the *inf-derivative* of that measure as

$$d^\wedge M(x) := \inf_{G: x \in G} M(G),$$

where G is allowed to vary across Borel sets. Here $d^\wedge M$ will be a lower semi-continuous function. Similarly, given a lower semi-continuous function f we can define the *inf-integral* of the function via

$$f^\wedge(G) := \inf_{x \in G} f(x).$$

Again, G here is any Borel set. It can be shown that f^\wedge is then an inf-measure. Hence one can develop the theory either in terms of inf-measures or in terms of lower semi-continuous functions.

A mathematically-sophisticated development of sup-measures, as well as a detailed account of connections with probability, optimization and analysis literature, is given in the excellent article [Ver97]. Unfortunately, Vervaat passed away before his work was published, and hence his work is only published somewhat obscurely.

The papers [RR91] and [RR92] extend and apply many of the ideas in [Ver97] to construct general random upper semi-continuous functions. These works are closely related to ours in that they construct a version of the process W_λ . However, they are strongly connected to the framework of choice optimization.

Various authors have addressed spatial effects in different contexts. For example, [dHL01] studies the probability that the maximum of some sequence of random functions $\xi_i(t)$ exceeds a deterministic function f , with $t \in [0, 1]$. This is similar to our framework in that they permit spatially-inhomogeneous shifts, but they do not study the distribution of the location of extremal events. Statistical estimation of spatial extremes in the context of max-stable processes has been used to study various geophysical processes [Smi90][DG12]. These works focus primarily on using spectral representations of max stable processes to tackle challenging spatial statistical problems.

1.2.1 Analogue with stochastic integrals

The discussion above suggests that for most of the basic developments of classical stochastic processes there is an analogue in terms of extreme values. In all cases the main difference is that taking sums of some underlying variables is replaced by taking the minima. In this light, one can see the theory of extreme values as a version of stochastic processes where the algebraic operator “+” is replaced with “min”. Extending this analogy, a family of algebraic operations similar to the ring $(+, \times)$, may be obtained by considering the operations $(\min, +)$. The resulting algebraic structure¹ $(\min, +)$ is known as a *tropical algebra*.

Our discussion of extremal processes above has not made use of “multiplication” of processes. However, this is critical for the construction of stochastic integrals. Recall that with classical Brownian motion one defines

$$\int_0^t H dB = \lim_{n \rightarrow \infty} \sum_{i=1}^n H_{t_{i-1}} (B_{t_i} - B_{t_{i-1}}),$$

where t_i describe partitions of $[0, t]$.

Since in studying extreme values we are replacing sums with mins and multiplication with addition, we thus ought to define our extremal stochastic integrals via

$$\text{“} \int_0^t g dW_\lambda \text{”} = \lim_{n \rightarrow \infty} \min_{i < nt} g \left(\frac{i}{n} \right) + n X_{\frac{i}{n}}.$$

In this paper we study these processes and give explicit formulae for the distribution of their first k mins and argmins.

¹Technically this is a semi-ring and not a proper ring, because the min operator is not invertible. However, this will not be important for our purposes.

1.3 Applications and extensions

There are several extensions and special cases of interest that can be derived from our main results. In this section we briefly describe a few promising examples.

1.3.1 Bayesian estimation via k -argmins

Suppose that the server problem in the introduction is repeated many times. That is, one sequentially sends requests to servers, and observes which were the first k servers to respond, as well as their response times. One can then rightfully ask: can we infer the functions g and λ ?

This question naturally fits a non-parametric *Bayesian* framework in which the functions g and λ are assumed to be unknown, and to be distributed according to a *prior* measure over functions. The explicit formulas developed in this paper would be nothing but the likelihood of the observations given the unknown functions g , λ . With these formulas at hand, a Markov chain Monte Carlo algorithm could be used to estimate expectations under the *posterior* distribution of unknowns given observations. In this way we could make predictions about where the location of the next fastest server will be and quantify how confident we are about our prediction.

1.3.2 Dependence of rates and delays on density

Suppose the points x_1, \dots, x_n are distributed according to the density $\rho(x)dx$. There may be applications where the rate function λ and the latency function g depend on x only through the density ρ of points around x . In particular, we may consider a model of the form

$$\lambda(x) := \Lambda(\rho(x)), \quad g(x) := G(\rho(x)),$$

where Λ and G are scalar functions. For example, we may imagine a parametric model of the form

$$\Lambda(t) = t^\alpha, \quad G(t) = t^\beta,$$

where α and β are real numbers. Estimating the functions λ and g then reduces to learning the parameters α and β . Estimating these parameters may provide valuable qualitative, as well as quantitative, information about the structure of latency and rate patterns due to high or low concentration of servers.

1.3.3 Extension to weakly-correlated response times

The response times $f_n(x_1), \dots, f_n(x_n)$ in Theorem 1.3 were assumed to be independent. Nevertheless, it is possible to extend our results to *weakly-correlated* processing times, allowing spatially correlated perturbations. Precisely, one could consider =

$$f_n(x_i) = \frac{n}{\lambda(x_i)} \xi_i + g(x_i)$$

where g is a random field independent from $\{\xi_1, \dots, \xi_n\}$ for which, with probability one,

$$\|g\|_\infty < \infty,$$

and where the ξ_i are independent. In this context, to obtain the asymptotic distribution of the first k argmins $x^{(1)}, \dots, x^{(k)}$ of f_n we may use the independence lemma to first obtain the asymptotic distribution of $(x^{(1)}, \dots, x^{(k)})|g$ using Corollary 1.4 and then integrate with respect to the distribution of g . The latency function g could be modeled according to a Gaussian random field.

1.3.4 Extensions to network structures

In the setting of server responses, it may be more accurate to consider a graph structure as opposed to an Euclidean one. Thus, consider a very large graph G_n with vertex set V_n . Suppose that to each node v in the graph we associate a server which can finish a task in time f_v . It is natural to ask whether one can learn the latency, rate, and weak correlation structure of processing times in such a large graph.

A possible practical approach to answer this question is to use an embedding of the set of nodes V_n into Euclidean space. For example, we can consider a spectral map $\Theta: V_n \rightarrow \mathbb{R}^d$ constructed using the graph Laplacian and then imagine that the servers are actually located at the “geographic” locations $\Theta(v_1), \dots, \Theta(v_n)$. One could then use the Bayesian inference methods described above to estimate the desired quantities. The idea of using the map Θ is that points v, \tilde{v} are “close” when the points $\Theta(v), \Theta(\tilde{v})$ are close to each other; in this way we translate the non-geographic information contained in G_n into geographic information which fits the set-up considered in this paper. Specific applications of these ideas are to be explored in the future.

1.4 Outline

The remainder of the paper is organized as follows. In Section 2 we collect some background results and we establish some properties of the process W_λ . The main results are proved in Section 3. We conclude in Section 4 with a short illustrative example.

2 Preliminaries

This section contains background results that will be employed in the proof of our main theorems. Subsection 2.1 describes the function space $\mathcal{S}(\overline{D})$ of lower semicontinuous functions endowed with a suitable topology. Subsection 2.2 establishes some properties and constructions of the process W_λ .

2.1 The space $\mathcal{S}(\overline{D})$ and the topology of Γ -convergence

Our goal here is to introduce the topology of Γ -convergence on the space of lower semicontinuous functions. This topology generates a notion of convergence which preserves the structure of minima and minimizers. Γ -convergence is also known as epi-convergence in the probability literature [Ver97]. This topology has found application in many fields, e.g. materials science, Ginzburg-Landau theory, and image processing. We will follow the presentation in [DM93], Chapter 10.

For any metric space X we define $\mathcal{S}(X)$ to be the family of lower semi-continuous functions on X taking values in $\mathbb{R} \cup \{\infty\}$. The following definition is standard:

Definition 2.1. *Let X be a metric space. A sequence $\{f_n\}_{n=1}^\infty \subset \mathcal{S}(X)$ is said to Γ -converge to $f \in \mathcal{S}(X)$ (written $f_n \xrightarrow{\Gamma} f$) if*

1. For all $x \in X$ there exists a sequence satisfying $x_n \rightarrow x$ such that

$$f(x) \geq \limsup_{n \rightarrow \infty} f_n(x_n). \quad (2.5)$$

2. For all $x_n \rightarrow x$ we have that

$$f(x) \leq \liminf_{n \rightarrow \infty} f_n(x_n). \quad (2.6)$$

In essence this definition requires that the limiting object f takes values below any possible limit, but that the value of f is obtained by some appropriately chosen “recovery sequence”. This notion of convergence imposes essentially the weakest conditions needed to guarantee convergence of minima and minimizers. This is manifest in the following proposition which is of high relevance for our purposes:

Proposition 2.2. *Suppose that $f_n \xrightarrow{\Gamma} f$, in some metric space X . Then $\min f_n \rightarrow \min f$. If X is also compact then if $x_n^* \in \operatorname{argmin} f_n$ then any limit point of x_n^* will be a minimizer of f .*

It is natural to seek a topology which describes Γ -convergence. To this end, we define the following topologies:

Definition 2.3. 1. We let σ^+ be the topology generated by sets of the form $\left\{ f \in \mathcal{S}(X) : \inf_U f < t \right\}$, where U is any open subset of X and $t \in \mathbb{R}$.

2. We let σ^- be the topology generated by sets of the form $\left\{ f \in \mathcal{S}(X) : \inf_K f > t \right\}$, where K is any compact subset of X and $t \in \mathbb{R}$.

3. We let σ be the smallest topology containing σ^+ and σ^- .

These topologies permit the measurement of minima on open and closed sets. It turns out that this topology is equivalent to Γ -convergence in the following sense (c.f. Theorem 10.17 in [DM93])

Proposition 2.4. *Let X be a metric space. A sequence $\{f_n\}_{n=1}^\infty \subset \mathcal{S}(X)$ converges in σ if and only if it Γ -converges.*

This space of functions is somewhat different from many of the standard function spaces. For example, it possesses the following compactness property (c.f. Theorem 10.6 in [DM93]):

Theorem 2.5. *The topological space $(\mathcal{S}(X), \sigma)$ is a compact space.*

An immediate application, which is of importance to the present investigation, is the following:

Corollary 2.6. *Any sequence $\{f_n\} \in \mathcal{S}(X)$ has a subsequence which Γ -converges.*

In general the topology generated by Γ -convergence will not be metrizable (or even Hausdorff). However, in the setting where X is compact we have the following as a consequence of Corollary 10.23 in [DM93]; see also Theorem 5.5 in [Ver97]:

Proposition 2.7. *Let X be a compact metric space. Then σ is metrizable.*

Finally, in this paper we will work with $X = \overline{D}$ where D is a bounded domain in \mathbb{R}^d . This allows us to characterize the topology in terms of cubes. The following is a direct application of Theorem 5.3 in [Ver97]:

Proposition 2.8. *For $X := \overline{D}$ where $D \subset \mathbb{R}^d$ is an open bounded domain, σ^\pm are generated by evaluating the infimum of functions over U being open cubes and K being closed cubes (as opposed to arbitrary open and closed sets).*

2.2 The process W_λ

In this subsection we establish some properties of the extreme value process W_λ introduced in Definition 1.1. We first show the existence and uniqueness of W_λ , and then the existence and uniqueness of its k -th argmins. We will denote by $\mathcal{P}(\mathcal{S})$ the space of probability measures over the space of lower semi-continuous functions on \bar{D} .

2.2.1 Existence and uniqueness of the process W_λ

We first show that the distribution of W_λ is a well-defined object.

Proposition 2.9. *Any two random variables with values in $\mathcal{S}(\bar{D})$ satisfying Definition 1.1 have the same distribution.*

Proof. Suppose that $\mathbb{P}_1, \mathbb{P}_2 \in \mathcal{P}(\mathcal{S})$ are the distributions of two random variables satisfying Definition 1.1. Let

$$\mathcal{A} = \{E \subset \mathcal{S}(D) \mid E \in \sigma, \mathbb{P}_1(E) = \mathbb{P}_2(E)\}.$$

It is straightforward to verify that \mathcal{A} is a λ -system.

Next, by Property 1 in Definition 1.1, we have that any set of the form $\{\inf_K f > s\}$ is in \mathcal{A} for any closed cube K and real number s . Similarly, by taking a limit of cubes from the outside and again using Property 1 in Definition 1.1 we have that any set of the form $\{\int_U f < t\}$ is in \mathcal{A} for any open cube U and real number t . By using Property 2 in Definition 1.1 we will have that any finite intersection of sets of these forms will also be in \mathcal{A} . Thus \mathcal{A} contains the π -system generated by sets of these forms. By the π - λ theorem, \mathcal{A} contains the sigma algebra generated by sets of these forms. By Proposition 2.8 it follows that $\mathcal{A} = \sigma$. This implies that $\mathbb{P}_1 = \mathbb{P}_2$, which concludes the proof. \square

Now we will demonstrate the existence of W_λ by constructing an appropriate approximating sequence.

Proposition 2.10. *Let $\{x_1, x_2, \dots\}$ be a countable subset of \bar{D} and for every $n \in \mathbb{N}$ let*

$$\mu_n := \frac{1}{n} \sum_{i=1}^n \delta_{x_i}.$$

We assume that μ_n converges weakly towards the distribution with density proportional to $\lambda(x)$. Let $\{\xi_1, \xi_2, \dots\}$ be i.i.d. exponentially-distributed random variables with rate one. Then, the random functions

$$W_n(x) := \begin{cases} \frac{n}{\bar{\lambda}} \xi_i, & \text{if } x = x_i, \\ +\infty, & \text{otherwise,} \end{cases}$$

converge weakly (in $\mathcal{P}(\mathcal{S})$) towards W_λ , where in the formula for W_n we are using

$$\bar{\lambda} := \int_D \lambda(x) dx.$$

Proof. Since the space $(\mathcal{S}(\bar{D}), \sigma)$ is compact, any sequence in $\mathcal{P}(\mathcal{S})$ is tight. Therefore, using that $(\mathcal{S}(\bar{D}), \sigma)$ is a separable, complete, metrizable space, Prokhorov's theorem implies that W_n must have a limit point (in the sense of weak convergence of measures) in $\mathcal{P}(\mathcal{S})$.

Let \tilde{W} be some limit point of the W_n . Our goal is to show that \tilde{W} satisfies the conditions given in the definition for W_λ . To that end, first consider a closed set C . From the definition of W_n , at every point x_i , $i=1\dots n$ we have an independent exponential variable, with rate $\frac{\lambda}{n}$. This then implies that $\min_{C \cap \bar{D}} W_n$ is exponentially distributed with rate $\mu_n(C \cap \bar{D})\lambda$. As $\mu_n \rightarrow \frac{\lambda}{\bar{\lambda}}$, we have that

$$\lim_{n \rightarrow \infty} \min_{C \cap \bar{D}} W_n$$

is exponentially distributed with rate $\int_{C \cap \bar{D}} \lambda(x) dx$. Since taking mins over closed sets is measurable in $\mathcal{P}(\mathcal{S})$, \tilde{W} must satisfy the first point in the definition of W_λ .

For the second property, we notice that the min over a finite number of disjoint closed sets will be independent under W_n , as the ξ_i are all independent. These events will also be in the topology σ . Hence by taking limits we obtain that \tilde{W} will satisfy the second property in the definition of W_λ . The uniqueness of W_λ given by Proposition 2.9 then gives that the unique limit point of W_n in $\mathcal{P}(\mathcal{S})$ is precisely W_λ , which completes the proof. \square

The previous construction of W_λ closely mirrors the framework described in the introduction: namely that we trace the rescaled minimum of many exponentially-distributed variables. An alternative construction (which can be connected to the Poisson process construction of extremal processes in [Pic71]) is also possible. This construction makes certain properties of the W_λ simpler to visualize, and we include it for completeness.

Proposition 2.11. *Let $\{x_1, x_2, \dots\}$ be a sequence of points, i.i.d according to the probability measure $\frac{\lambda}{\bar{\lambda}} dx$. Let $\{\xi_1, \xi_2, \dots\}$ be an i.i.d. sequence of random variables distributed as $\exp(\bar{\lambda})$. Define the following random functions:*

$$W_n(x) = \begin{cases} \sum_{j=1}^i \xi_j, & \text{for } x = x_i, \quad i \leq n, \\ +\infty, & \text{otherwise.} \end{cases}$$

Then the W_n converge weakly (in $\mathcal{P}(\mathcal{S})$) towards W_λ .

Proof. As in the previous construction, we only need to prove that W_n asymptotically satisfies Properties 1 and 2 from Definition 1.1. We will do this by direct computation.

For Property 1 in Definition 1.1, given any closed set C , the probability that x_k lies in C is equal to $\bar{\lambda}^{-1} \int_{C \cap \bar{D}} \lambda(x) dx =: \tilde{\lambda}_C$. Thus the first i such that $x_i \in C$ is geometrically distributed with parameter $\tilde{\lambda}_C$. Furthermore, $\sum_{j=1}^i \xi_j$ has an Erlang distribution with density $\frac{\bar{\lambda}^i t^{i-1} e^{-\bar{\lambda}t}}{(i-1)!}$. In turn

$$\mathbb{P} \left(\min_{x \in C \cap \bar{D}} W_n \geq r \right) = \int_r^\infty \sum_{j=1}^n (1 - \tilde{\lambda}_C)^{j-1} \tilde{\lambda}_C \frac{\bar{\lambda}^j s^{j-1} e^{-\bar{\lambda}s}}{(j-1)!} ds + \sum_{j=n+1}^\infty (1 - \tilde{\lambda}_C)^{j-1} \tilde{\lambda}_C,$$

where on the right we are using the fact that the choice of points x_i is independent of the values of the ξ_j . Taking a limit as $n \rightarrow \infty$ and simplifying the series, we find that

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\min_{x \in C \cap \bar{D}} W_n \geq r \right) = \int_r^\infty \tilde{\lambda}_C \bar{\lambda} \exp(-\tilde{\lambda}_C \bar{\lambda} s) ds.$$

This proves Property 1.

With regards to Property 2, for $r < s$ we can compute (letting p_n be the associated density function):

$$\begin{aligned}
p_n(\min_{C_1} W_n = r, \min_{C_2} W_n = s) &= \sum_{k_1=1}^n (1 - \tilde{\lambda}_{C_1} - \tilde{\lambda}_{C_2})^{k_1-1} \tilde{\lambda}_{C_1} \frac{\bar{\lambda}^{k_1} r^{k_1-1} e^{-\bar{\lambda}r}}{(k_1-1)!} \\
&\times \sum_{k_2=k_1+1}^n (1 - \tilde{\lambda}_{C_2})^{k_2-k_1-1} \tilde{\lambda}_{C_2} \frac{\tilde{\lambda}^{k_2-k_1} (s-r)^{k_2-k_1-1} e^{-\tilde{\lambda}(s-r)}}{(k_2-k_1-1)!} \\
&= \sum_{k_1=1}^n (1 - \tilde{\lambda}_{C_1} - \tilde{\lambda}_{C_2})^{k_1-1} \tilde{\lambda}_{C_1} \frac{\bar{\lambda}^{k_1} r^{k_1-1} e^{-\bar{\lambda}r}}{(k_1-1)!} \\
&\times \sum_{\tilde{k}=1}^{n-k_1} (1 - \tilde{\lambda}_{C_2})^{\tilde{k}-1} \tilde{\lambda}_{C_2} \frac{\tilde{\lambda}^{\tilde{k}} (s-r)^{\tilde{k}-1} e^{-\tilde{\lambda}(s-r)}}{(\tilde{k}-1)!}.
\end{aligned}$$

Taking $n \rightarrow \infty$ we then obtain

$$\begin{aligned}
&\lim_{n \rightarrow \infty} p_n(\min_{C_1} W_n = r, \min_{C_2} W_n = s) \\
&= \tilde{\lambda}_{C_1} \bar{\lambda} \exp(-\bar{\lambda}r + \bar{\lambda}r(1 - \tilde{\lambda}_{C_1} - \tilde{\lambda}_{C_2})) \times \tilde{\lambda}_{C_2} \bar{\lambda} \exp(-\tilde{\lambda}(s-r) + (s-r)\tilde{\lambda}(1 - \tilde{\lambda}_{C_2})) \\
&= \tilde{\lambda}_{C_1} \bar{\lambda} \exp(-\bar{\lambda}\tilde{\lambda}_{C_1}r) \times \tilde{\lambda}_{C_2} \bar{\lambda} \exp(-\tilde{\lambda}\tilde{\lambda}_{C_2}s).
\end{aligned}$$

This proves Property 2 in the case of two sets. The case with more than two sets is completely analogous. \square

Remark 2.12. *The previous construction highlights the memorylessness of the process W_λ . That is, the distribution of the k -th min can be found by restarting the process after the arrival of the $k-1$ -th min. This memorylessness will be very convenient in establishing integral formulas in Section 3. However, one cannot expect this property to hold for other extreme value distributions, see Remark 1.5.*

Remark 2.13. *There are several classical constructions of Brownian motion. Some, such as Lévy's piecewise linear construction, are easy to visualize but only converge as a measure on the space of continuous functions. Others, such as the "wavelet-type" construction of Lévy-Ciesielski, or the Fourier construction of Weiner, converge uniformly towards Brownian motion (see e.g. Chapter 3 in [SP14] for more details). Here we only characterize the convergence of the distribution of W_n towards W_λ as measures on the space of lower semi-continuous functions. One could seek to demonstrate that certain alternative constructions converges uniformly, in an appropriate metric on the space of lower semi-continuous functions as in Proposition 2.7 (see also [DM93], Proposition 10.21). This will be the subject of future analysis.*

2.2.2 Existence and uniqueness of k -th argmins of W_λ

Next we turn to studying the existence of k -th argmins for the process W_λ . We begin with a definition:

Definition 2.14. *Let $f \in \mathcal{S}(\bar{D})$. We define the minimum value of f*

$$m_1(f) := \min_{x \in \bar{D}} f(x),$$

as well as the set of minimizers of f

$$M_1(f) := \{x \in \overline{D} : f(x) = m_1(f)\}.$$

Moreover, having defined the numbers $m_1(f), \dots, m_{k-1}(f)$ and the sets $M_1(f), \dots, M_{k-1}(f)$ we define

$$m_k(f) := \inf_{x \in \overline{D} \setminus \bigcup_{i=1}^{k-1} M_i(f)} f(x) \quad (2.7)$$

and

$$M_k(f) := \{x \in \overline{D} : m_{k-1}(f) < f(x) = m_k(f)\}.$$

We refer to the elements of $M_k(f)$ as k -th argmins of f .

Remark 2.15. The set $M_1(f)$ is always non-empty due to the compactness of \overline{D} and the lower semi-continuity of f . The lower semi-continuity of f also implies that for every $k \in \mathbb{N}$, the set

$$\bigcup_{i=1}^k M_i(f)$$

is a closed set.

Remark 2.16. In general the sets $M_k(f)$ for $k > 1$ may be empty. For example, it should be clear that a continuous function f does not have 2-th argmins. In this paper, however, the random object that we consider is almost surely a lower semi-continuous function with well-defined (and unique) k -argmins for all k (see Proposition 2.19 below).

Remark 2.17. It is an easy exercise to show that if $M_k(f)$ is a non-empty set then $M_{k-1}(f)$ is also non-empty.

The following lemma will be critical in proving the existence of k -argmins of $W_\lambda + g$.

Lemma 2.18. Let $\{f_n\} \subset \mathcal{S}(\overline{D})$, let $k \in \mathbb{N}$ and suppose that:

1. $f_n \xrightarrow{\Gamma} f \in \mathcal{S}$.
2. The set $M_i(f)$ is a singleton for all $i \leq k$.
3. There exists a number $\beta > 0$ such that

$$d(M_i(f_n), M_j(f_n)) > \beta, \quad \forall n \in \mathbb{N}, \quad \forall i, j = 1, \dots, k+1, \quad i \neq j. \quad (2.8)$$

In the above, $d(M_i(f_n), M_j(f_n))$ is defined as

$$d(M_i(f_n), M_j(f_n)) := \inf\{|x - y| : x \in M_i(f_n), \quad y \in M_j(f_n)\}.$$

4. For n sufficiently large $M_i(f_n) \neq \emptyset$, for all $i \leq k+1$.

Then all limit points of $\{x_n^{k+1}\}$, where $x_n^{k+1} \in M_{k+1}(f_n)$, belong to $M_{k+1}(f)$. In particular, $M_{k+1}(f)$ is non-empty.

Proof. Step 1: We start by proving the case $k=1$ in order to illustrate the ideas.

Let $x^1 \in M_1(f)$ be the unique minimizer of f . By compactness of \overline{D} we know that $\{x_n^2\}_{n \in \mathbb{N}}$ converges up to subsequence towards a point $\tilde{x}^2 \in \overline{D}$; without loss of generality we assume that the full sequence converges towards \tilde{x}^2 . We need to show that $\tilde{x}^2 \in M_2(f)$.

First, we claim that $\tilde{x}^2 \neq x^1$, and so $\tilde{x}^2 \notin M_1(f)$. To see this, let x_n^1 be a minimizer for f_n . By the Γ -convergence assumption and the fact that x^1 is the unique minimizer of f , it follows that $x_n^1 \rightarrow x^1$. In particular, for large enough n we have

$$|x_n^1 - x^1| < \beta/2,$$

where β is as in (2.8). From the triangle inequality it follows that for large enough n

$$|x_n^2 - x^1| \geq |x_n^1 - x_n^2| - |x_n^1 - x^1| \geq \beta - \beta/2 = \beta/2.$$

Therefore,

$$|\tilde{x}^2 - x^1| \geq \beta/2 > 0,$$

establishing the claim.

Next, we show that for arbitrary $x \in \overline{D} \setminus M_1(f) = \overline{D} \setminus \{x^1\}$ we have that $f(\tilde{x}^2) \leq f(x)$. This implies that $\tilde{x}^2 \in M_1(f) \cup M_2(f)$, which together with the above shows that $\tilde{x}^2 \in M_2(f)$. Using the lim-sup inequality we can find a sequence $\{x_n\}_{n \in \mathbb{N}}$ converging towards x for which

$$\limsup_{n \rightarrow \infty} f_n(x_n) \leq f(x).$$

Notice that for all large enough n , $x_n \neq x_n^1$. Indeed, if that was not the case, the limit of $\{x_n\}_{n \in \mathbb{N}}$ would be x^1 which contradicts that $x \neq x^1$. In particular, the value of $f_n(x_n^2) \leq f_n(x_n)$. Hence,

$$f(\tilde{x}^2) \leq \liminf_{n \rightarrow \infty} f_n(x_n^2) \leq \limsup_{n \rightarrow \infty} f_n(x_n) \leq f(x),$$

where the first inequality follows from the liminf inequality given that $x_n^2 \rightarrow \tilde{x}^2$.

Step 2: Generalizing to arbitrary k follows the same ideas.

Let, for $i=1, \dots, k$, $x^i \in M_i(f)$ be the unique i -th argmin of f and let, for $i=1, \dots, k+1$, $x_n^i \in M_i(f_n)$. We assume without loss of generality that $x_n^i \rightarrow \tilde{x}^i$ for some $\tilde{x}^i \in \overline{D}$. Our goal is to prove that $\tilde{x}^{k+1} \in M_{k+1}(f)$.

Restricted to the set $B(\tilde{x}^i, \beta/2)$ we have that for n sufficiently large x_n^i minimizes f_n . Γ -convergence then implies that $f_n(x_n^i) \rightarrow f(\tilde{x}^i)$. By the definition of $M_i(f_n)$ we have that, for $i < j$, $f_n(x_n^i) < f_n(x_n^j)$ and hence $f(\tilde{x}^i) \leq f(\tilde{x}^j)$. Note that $\tilde{x}^1 = x^1$, simply by Γ -convergence (on the whole space \overline{D}) and the fact that $M_1(f)$ is a singleton.

Now, we claim that for all $i=1, \dots, k$ we have that $\tilde{x}^i = x^i$. We proceed by induction. The base case was already proved. Suppose that $\tilde{x}^i = x^i$ for all $i=1, \dots, j < k$. If $\tilde{x}^{j+1} \neq x^{j+1}$ then we must have (by the fact that $M_i(f)$ is a singleton) that $f(\tilde{x}^j) < f(x^{j+1}) < f(\tilde{x}^{j+1})$. By Γ -convergence there exists a sequence $x_n \rightarrow x^{j+1}$ so that $f_n(x_n) \rightarrow f(x^{j+1})$. This then implies that $f_n(x_n^j) < f_n(x_n) < f_n(x_n^{j+1})$ for n sufficiently large. This violates the fact that $x_n^{j+1} \in M_{j+1}(f_n)$, which then proves the claim.

Finally, let $x \in \overline{D} \setminus \bigcup_{i=1}^k M_i(f)$ be arbitrary. We will show that $f(\tilde{x}^{k+1}) \leq f(x)$ which implies that $\tilde{x}^{k+1} \in M_{k+1}(f)$. By Γ -convergence, there exists a sequence $x_n \rightarrow x$ with $f_n(x_n) \rightarrow f(x)$. Since $x \notin \bigcup_{i=1}^k M_i(f)$ we have that, for $i=1, \dots, k$, $f(x) > f(x^i)$ and so, for n large enough, $f_n(x_n) > f_n(x_n^i)$. Because $x_n^{k+1} \in M_{k+1}(f_n)$, this in turn implies that $f_n(x_n) \geq f_n(x_n^{k+1})$. Taking limits concludes the proof. \square

We are now ready to show the existence and uniqueness of k -argmins for $W_\lambda + g$.

Proposition 2.19. *Suppose that $g: \overline{D} \rightarrow \mathbb{R}$ is a continuous function. Then with probability one, for every $k \in \mathbb{N}$, $M_k(W_\lambda + g)$ is a singleton. In other words $W_\lambda + g$ has well-defined and unique k -th argmins for every $k \in \mathbb{N}$.*

Proof. By Proposition 2.11 we can assume without loss of generality that, with probability one, W_n as defined in (2.11) Γ -converges towards W_λ . Since g is continuous it follows that, with probability one,

$$W_n + g \xrightarrow{\Gamma} W_\lambda + g.$$

From the construction of the W_n it is easy to see that, with probability one, the sequence of functions $\{W_n + g\}_{n \in \mathbb{N}}$ satisfies conditions 1-3 in Lemma 2.18. In order to show the existence of k -argmins it is enough to show that with probability one $M_k(W_\lambda + g)$ cannot have more than two elements (proving in this way existence and uniqueness).

We only consider the case $k=1$ as the general case is proved similarly. Let \mathcal{F}_l be the family of dyadic cubes Q with side length 2^{-l} . Suppose that $M_1(W_\lambda + g)$ has at least two elements. Then for some $l \in \mathbb{N}$ we have that $\min_Q W_\lambda + g = \min_{Q'} W_\lambda + g$ for two disjoint $Q, Q' \in \mathcal{F}_l$. However, by Properties 1 and 2 of Definition 1.1 we have that

$$\mathbb{P} \left(\bigcup_{l=1}^{\infty} \bigcup_{\substack{Q, Q' \in \mathcal{F}_l \\ Q \cap Q' = \emptyset}} \left\{ \omega \in \Omega : \min_{x \in Q} W + g = \min_{x \in Q'} W + g \right\} \right) = 0.$$

This then implies that $M_1(W_\lambda + g)$ must be a singleton (with probability 1). □

Remark 2.20. *We have used the fact that argmins are locally isolated in a very strong way. Ideas of this type have previously been applied in the context of Cahn-Hilliard phase transitions, see e.g. [KS89].*

3 Proof of the main results

This section contains the proofs of the main results of the paper, Theorems 1.2 and 1.3

3.1 Proof of Theorem 1.2

We already showed in Proposition 2.19 that $W_\lambda + g$ has well-defined k -th argmins. Our aim now is to establish equation 1.2, that gives a formula for their joint distribution. We will consider first the case $k=1$ in Lemma 3.1, and then extend the proof to the case $k > 1$.

Before delving into the proofs we recall some notation —see equation (1.3). We denote by μ_λ the probability measure on \overline{D} with density proportional to λ ,

$$\frac{d\mu_\lambda(x)}{dx} \propto \lambda(x), \quad x \in \overline{D}.$$

The normalizing constant is denoted by $\bar{\lambda}$, that is,

$$\bar{\lambda} := \int_D \lambda(x) dx.$$

We also recall the definition of the cumulative distribution function

$$H_{\lambda,g}(s) := \mu_\lambda(\{x \in D : g(x) \leq s\}).$$

Henceforth we will omit the dependence of μ_λ and $H_{\lambda,g}$ on λ and g from the notation, and write simply μ and H .

Lemma 3.1. *Let $g \in \mathcal{S}(\overline{D})$, and let $X^{(1)}$, $\tau^{(1)}$ be random variables representing the global minimizer and minimum value of $W_\lambda + g$. Then the joint density of these random variables is given by*

$$\frac{\mathbb{P}(X^{(1)} \in dx, \tau^{(1)} \in d\tau)}{dx d\tau} = \Psi(x, \tau, g) = \chi_{\tau > g(x)} \exp\left(-\bar{\lambda} \int_{-\infty}^{\tau} H(s) ds\right). \quad (3.9)$$

In turn the density of $X^{(1)}$ is given by

$$\frac{\mathbb{P}(X^{(1)} \in dx)}{dx} = \Phi(x, g) = \int_{g(x)}^{\infty} \exp\left(-\bar{\lambda} \int_{-\infty}^t H(s) ds\right) dt. \quad (3.10)$$

Proof. Let $\{x_1, x_2, \dots\}$ be a countable set of points in D that we use to construct discrete approximations of μ and H . Precisely, we set, for every $n \in \mathbb{N}$,

$$\mu_n := \frac{1}{n} \sum_{i=1}^n \delta_{x_i}, \quad H_n(\alpha) := \mu_n(\{x \in D : g(x) \leq \alpha\}).$$

The points $\{x_1, x_2, \dots\}$ can be chosen so that μ_n converges weakly towards μ , and H_n converges uniformly towards H .

Let $\{\xi_1, \xi_2, \dots\}$ be a sequence of i.i.d. exponential random variables with rate one. Fix $n \in \mathbb{N}$ and consider the (random) function

$$\bar{f}_n(x) := \begin{cases} \frac{n}{\bar{\lambda}} \xi_i + g(x_i), & \text{if } x = x_i, \quad i = 1, \dots, n, \\ +\infty, & \text{otherwise.} \end{cases}$$

Let $X_n^{(1)}$ and $\tau_n^{(1)}$ denote the minimizer and minimum of the function \bar{f}_n . We may compute, for $r > 0$,

$$\begin{aligned} \mathbb{P}\left(X_n^{(1)} = x_j, \tau_n^{(1)} = r + g(x_j)\right) &= \mathbb{P}\left(\frac{n}{\bar{\lambda}} \xi_j = r, \frac{n}{\bar{\lambda}} \xi_j \leq \frac{n}{\bar{\lambda}} \xi_i + (g(x_i) - g(x_j)) \forall i\right) \\ &= \mathbb{P}\left(\xi_j = r \frac{\bar{\lambda}}{n}, \xi_i > \frac{\bar{\lambda}}{n} (r + g(x_j) - g(x_i))\right) \\ &= \exp\left(-r \frac{\bar{\lambda}}{n}\right) \int_{\frac{\bar{\lambda}}{n}(r - (g(x_1) - g(x_j)))_+}^{\bar{\lambda}} \cdots \int_{\frac{\bar{\lambda}}{n}(r - (g(x_n) - g(x_j)))_+}^{\bar{\lambda}} \prod_{i \neq j} \exp(-r_i) dr_i \\ &= \exp\left(-r \frac{\bar{\lambda}}{n}\right) \exp\left(\frac{\bar{\lambda}}{n} \sum_{g(x_i) < r + g(x_j)} (g(x_i) - g(x_j) - r)\right). \end{aligned}$$

Using the layer cake decomposition we may write

$$\mathbb{P}(X_n^{(1)} = x_j, \tau_n^{(1)} = r + g(x_j)) = \exp\left(-r \frac{\bar{\lambda}}{n}\right) \exp\left(-\bar{\lambda} \int_{-\infty}^{r + g(x_j)} H_n(s) ds\right).$$

Replacing $r + g(x_j)$ with $\tau_n^{(1)}$ and letting $n \rightarrow \infty$ we obtain

$$\frac{\mathbb{P}(X^{(1)} \in dx, \tau^{(1)} \in d\tau)}{dx d\tau} = \Psi(x, \tau, g) := \chi_{\tau > g(x)} \exp\left(-\bar{\lambda} \int_{-\infty}^{\tau} H(s) ds\right). \quad (3.11)$$

Now taking the integral in τ gives

$$\frac{\mathbb{P}(X^{(1)} \in dx)}{dx} = \Phi(x, g) := \int_{g(x)}^{\infty} \exp\left(-\bar{\lambda} \int_{-\infty}^t H(s) ds\right) dt.$$

Proposition 2.10 implies that, with probability one, $f_n \xrightarrow{\Gamma} W_\lambda + g$. This in turn implies the desired result thanks to Proposition 2.2. \square

We are now ready to establish the formulas for the joint distribution of the first k -th argmins.

Proof of Theorem 1.2: We complete the proof of Theorem 1.2 by obtaining the joint distribution of $(X^{(1)}, \dots, X^{(k)})$.

We first observe that

$$\begin{aligned} \mathbb{P}(X^{(i)} \in A | X^{(1)}, \dots, X^{(i-1)}, \tau^{(1)}, \dots, \tau^{(i-1)}) &= \mathbb{P}(X^{(i)} \in A | \tau^{(i-1)}), \\ \mathbb{P}(X^{(i)} \in A, \tau^{(i)} \in B | X^{(1)}, \dots, X^{(i-1)}, \tau^{(1)}, \dots, \tau^{(i-1)}) &= \mathbb{P}(X^{(i)} \in A, \tau^{(i)} \in B | \tau^{(i-1)}). \end{aligned}$$

Bayes' rule then implies that

$$\frac{\mathbb{P}(X^{(1)} \in dx_1, \dots, X^{(k)} \in dx_k)}{dx_1 \dots dx_k} = \int_{\mathbb{R}^{k-1}} \frac{\mathbb{P}(X^{(k)} | \tau^{(k-1)})}{dx_k} \frac{\mathbb{P}(X^{(1)}, \tau^{(1)})}{dx_1 dr_1} \prod_{j=2}^{k-1} \frac{\mathbb{P}(X^{(j)}, \tau^{(j)} | \tau^{(j-1)})}{dx_j dr_j} dr_1 \dots dr_{k-1}. \quad (3.12)$$

Here we use r_i as a dummy variable to avoid overloading notation for τ . Using the fact that exponentials are memoryless (along with the construction in Proposition 2.10 and the convergence in the previous proof) we then obtain

$$\frac{\mathbb{P}(X^{(1)} \in dx_1, \dots, X^{(k)} \in dx_k)}{dx_1 \dots dx_k} = \int_{\mathbb{R}^{k-1}} \Phi(x_k, (g - r_{k-1})^+) \Psi(x_1, r_1, g) \quad (3.13)$$

$$\times \prod_{j=2}^{k-1} \Psi(x_j, r_j - r_{j-1}, (g - r_{j-1})^+) dr_1 \dots dr_{k-1}. \quad (3.14)$$

This matches the formula in Theorem 1.2 and completes the proof. \square

3.2 Proof of Theorem 1.3

Proof of Theorem 1.3. To begin with, we provide the (classical) proof that

$$n \min\{\xi_1 \dots \xi_n\} \rightarrow^d \exp(1).$$

Define $Y_n := n \min\{\xi_1 \dots \xi_n\}$. Then

$$\mathbb{P}(Y_n > t) = \mathbb{P}(\xi_1 > t/n, \dots, \xi_n > t/n).$$

By independence

$$\mathbb{P}(Y_n > t) = \prod_{j=1}^n (1 - F_j(t/n)).$$

Therefore, using hypothesis (1.4),

$$\lim_{n \rightarrow \infty} \mathbb{P}(Y_n > t) = e^{-t},$$

proving the claim.

We then remark that the exact same proof holds if we only consider subsets of the sample points (e.g. $\{x_i \in C\}$). This, along with independence of the ξ_i implies that, for f_n as in Theorem 1.3, we have that $nf_n \rightarrow W_\lambda + g$, where this convergence is in $\mathcal{P}(\mathcal{S})$.

The formula for the limiting joint distribution of the argmins then follows as in the proof of Lemma 3.1. □

Remark 3.2. *We remark that convergence in $\mathcal{P}(\mathcal{S})$ does not necessarily imply convergence in distribution of the sets M_k (except for the case $k=1$). This is because small scale features may be lost in the limiting procedure at the level of k -argmins. However, the “pointwise” construction that we use here preserves the convergence of k -argmins, since all argmins are locally isolated.*

4 Computational example on $D = (0, 1)$

We illustrate some of our results by considering a family of distributions obtained from computing our “stochastic integrals” on the unit interval $[0, 1]$. Consider the function

$$g(x) := x^2.$$

Obviously the distribution of the minimizers of g is the Dirac delta measure at zero. Let us now consider the random function

$$x \mapsto \delta W_\lambda(x) + g(x),$$

for some parameter $\delta > 0$, and for $\lambda \equiv 1$. The distribution of the minimizers of $\delta W_\lambda + g$ will naturally depend on the value of δ and is expected to be close to a Dirac delta measure for small values of δ , whereas is expected to be essentially uniform for large values of δ . We will then interpret the distributions of minimizers of $\delta W_\lambda + g$ as regularizations of the Dirac delta at zero.

We can compute explicitly the densities for the distribution of minimizers using (1.2). Indeed the distribution of the minimizer of $\delta W_\lambda + g$ is given by

$$\rho_\delta(y) := \frac{2}{\delta} \int_y^1 x \exp\left(-\frac{2x^3}{3\delta}\right) dx + \exp\left(-\frac{2\delta}{3}\right), \quad y \in [0, 1].$$

Figure 1 plots this density for several representative values of δ .

Our results give us information of *where* the minimizer is located. In particular it allows us to quantify the likelihood to find a minimizer of the process in any given subregion of D .

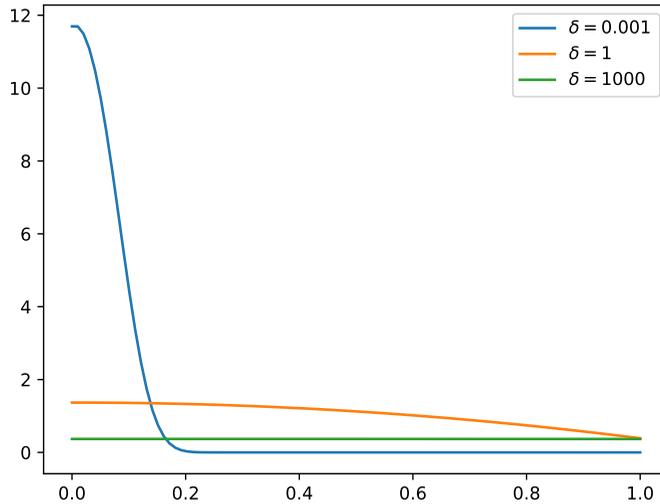


Figure 1: The distribution of minimizers of $x^2 + \delta W_\lambda$, for different values of δ .

References

- [DG12] A. C. Davison and M. M. Gholamrezaee. Geostatistics of extremes. *Proc. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.*, 468(2138):581–608, 2012.
- [dH84] Laurens de Haan. A spectral representation for max-stable processes. *Ann. Probab.*, 12(4):1194–1204, 1984.
- [dHL01] Laurens de Haan and Tao Lin. On convergence toward an extreme value distribution in $C[0,1]$. *Ann. Probab.*, 29(1):467–483, 2001.
- [dHP84] Laurens de Haan and James Pickands, III. Stationary min-stable stochastic processes. In *Statistical extremes and applications (Vimeiro, 1983)*, volume 131 of *NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci.*, pages 483–489. Reidel, Dordrecht, 1984.
- [DM93] Gianni Dal Maso. *An Introduction to Γ -convergence*. Springer, 1993.
- [Dwa64] Meyer Dwass. Extremal processes. *Ann. Math. Statist*, 35:1718–1725, 1964.
- [Gne43] Boris Gnedenko. Sur la distribution limite du terme maximum d’une série aléatoire. *Ann. of Math. (2)*, 44:423–453, 1943.
- [KS89] Robert V. Kohn and Peter Sternberg. Local minimisers and singular perturbations. *Proc. Roy. Soc. Edinburgh Sect. A*, 111(1-2):69–84, 1989.
- [Lam64] John Lamperti. On extreme order statistics. *Ann. Math. Statist*, 35:1726–1737, 1964.

- [LLR83] M. R. Leadbetter, Georg Lindgren, and Holger Rootzén. *Extremes and related properties of random sequences and processes*. Springer Series in Statistics. Springer-Verlag, New York-Berlin, 1983.
- [Pic71] James Pickands, III. The two-dimensional Poisson process and extremal processes. *J. Appl. Probability*, 8:745–756, 1971.
- [Res13] Sidney I Resnick. *Extreme values, regular variation and point processes*. Springer, 2013.
- [RR91] Sidney I. Resnick and Rishin Roy. Random usc functions, max-stable processes and continuous choice. *Ann. Appl. Probab.*, 1(2):267–292, 1991.
- [RR92] Sidney I. Resnick and Rishin Roy. On min-stable horse races with infinitely many horses. *Math. Social Sci.*, 23(1):119–145, 1992.
- [Smi90] Richard Smith. Max-stable processes and spatial extremes. *Unpublished manuscript*, 1990.
- [SP14] René L. Schilling and Lothar Partzsch. *Brownian motion: an Introduction to stochastic processes*. De Gruyter Graduate. De Gruyter, Berlin, second edition, 2014. An introduction to stochastic processes, With a chapter on simulation by Björn Böttcher.
- [Ver97] Wim Vervaat. Random upper semicontinuous functions and extremal processes. In *Probability and lattices*, volume 110 of *CWI Tract*, pages 1–56. Math. Centrum, Centrum Wisk. Inform., Amsterdam, 1997.