



Stochastic Systems

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Efficiency of Parallel and Restart Exploration Strategies in Model-Free Stochastic Simulations

Ernesto Garcia, Paola Bermolen, Matthieu Jonckheere, Seva Shneer

To cite this article:

Ernesto Garcia, Paola Bermolen, Matthieu Jonckheere, Seva Shneer (2026) Efficiency of Parallel and Restart Exploration Strategies in Model-Free Stochastic Simulations. *Stochastic Systems*

Published online in Articles in Advance 15 Jun 2026

. <https://doi.org/10.1287/stsy.2025.0108>

This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*Stochastic Systems*. Copyright © 2026 The Author(s). <https://doi.org/10.1287/stsy.2025.0108>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Copyright © 2026 The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>


Efficiency of Parallel and Restart Exploration Strategies in Model-Free Stochastic Simulations

Ernesto Garcia,^{a,*} Paola Bermolen,^b Matthieu Jonckheere,^a Seva Shneer^c

^aLAAS-CNRS, 31031 Toulouse, France; ^bFacultad de Ingeniería, CICADA, Universidad de la República, Montevideo 11200, Uruguay;

^cHeriot-Watt University, Edinburgh EH14 4AS, United Kingdom

*Corresponding author

Contact: egarciacig@laas.fr,  <https://orcid.org/0009-0007-3200-4830> (EG); paola@fing.edu.uy (PB); matthieu.jonckheere@laas.fr (MJ); v.shneer@hw.ac.uk (SS)

Received: March 26, 2025

Revised: December 22, 2025; May 7, 2026

Accepted: May 12, 2026

Published Online in Articles in Advance:
June 15, 2026

<https://doi.org/10.1287/stsy.2025.0108>

Copyright: © 2026 The Author(s)

Abstract. We analyze the efficiency of parallelization and restart mechanisms for stochastic simulations in model-free settings, where the underlying system dynamics are unknown. Such settings are common in Reinforcement Learning (RL) and rare-event estimation, where standard variance-reduction techniques like importance sampling are inapplicable. Focusing on the challenge of reaching rare states under a finite computational budget, we model exploration via random walks and Lévy processes. Based on rigorous probability analysis, our work reveals a phase transition in the success probability as a function of the number of parallel simulations: an optimal number N^* exists, balancing exploration diversity and time allocation per simulation. Beyond this threshold, performance degrades exponentially. Furthermore, we demonstrate that a restart strategy, which reallocates resources from stagnant trajectories to promising regions, can yield an exponential improvement in success probability. In the context of RL, these strategies can improve policy gradient methods by enabling more efficient state-space exploration, leading to more accurate policy gradient estimates.



Open Access Statement: This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “Stochastic Systems. Copyright © 2026 The Author(s). <https://doi.org/10.1287/stsy.2025.0108>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Funding: This research was supported by the SticAmsud LAGOON project, ANR EPLER, IRL-CNRS IFUMI-2030, and Action international CNRS.

Keywords: stochastic exploration • random walk • Lévy process • Reinforcement Learning • rare rewards

1. Introduction

The efficient simulation of stochastic systems is a cornerstone of many critical applications, ranging from rare-event estimation to gradient estimation in Reinforcement Learning (RL). In many practical scenarios, the underlying dynamics of a system are complex, poorly characterized, or entirely unknown, placing us essentially in a model-free setting. In such contexts, strategies for exploring the state space where the stochastic dynamics live and gathering informative samples are paramount. This work investigates two powerful and general mechanisms for enhancing the efficiency of stochastic simulations: parallelization and restarting. We aim to provide a rigorous probabilistic analysis of these strategies, with applications focused on key model-free problems such as estimating small probabilities for stochastic processes with unknown dynamics and the challenge of efficient exploration in RL.

A canonical problem in probability and statistics is the estimation of the probability that a stochastic process reaches a rare, distant state within a finite time horizon. Standard Monte Carlo methods fail as the event of interest is rarely observed. A common technique for accelerating rare-event simulation is importance sampling (Asmussen and Glynn 2007), which involves simulating the process under an alternative “twisted” measure that makes the rare event more likely, and then correcting this change of measure via likelihood ratios. However, importance sampling is not a viable solution in the model-free settings as its application requires exact knowledge of the underlying process dynamics to construct the optimal change of measure. When the dynamics are unknown or too complex to model—as is often the case in large-scale RL environments or with intractable stochastic processes—this precise knowledge is unavailable, rendering importance sampling inapplicable. Moreover, an estimation scheme for the change of measure is usually of very large variance. It is precisely this limitation that motivates the use of blind or model-free strategies, such as parallelization and restarting, which do not require explicit knowledge of the system dynamics.

Parallelization—running multiple independent simulations concurrently—seems a natural remedy, as it increases the chances of observing the rare event. However, under a fixed total computational budget, a critical trade-off emerges: should one allocate the entire budget to a single, long simulation, or distribute it among many shorter, parallel runs? The answer is not trivial, as splitting the budget too finely may deprive each simulation of the necessary time to reach the rare state. We analyze this trade-off quantitatively, identifying a sharp phase transition in the success probability as a function of the number of parallel runs.

Beyond parallelization, restart strategies offer another avenue for improvement. Instead of letting simulations run in potentially unfruitful regions of the state space, a restart mechanism halts them and reinitializes them from more promising states, effectively redirecting computational effort. This restarting mechanism is particularly powerful in scenarios where certain regions are more likely to lead to the target event. A well-designed restart policy can prevent the waste of resources on trajectories with a low probability of success and serves as a practical, model-free alternative to importance sampling.

Several recent works have addressed the efficient exploration challenge, in particular for RL. For example, the Go-Explore approach in Ecoffet et al. (2021) introduces an innovative solution by distinctly separating the exploration and exploitation phases. This method allows the algorithm to maintain a repository of promising states, which can be revisited for further exploration, significantly enhancing the ability to discover and leverage rare or hard-to-reach states. Another example is found in Mastropietro et al. (2025), based on the paradigm of the Fleming-Viot particle system, introduced by Burdzy et al. (1996). In this framework, a population of particles (which can represent different simulations in parallel) evolves over time, exploring the state space in parallel. When a particle falls into a less promising region, it is “restarted” by being replaced with a copy of another, more promising particle. This ensures that resources are concentrated on exploring the most fruitful areas of the state space.

Beyond RL, stochastic restarting has been studied extensively across disciplines such as statistical physics, computer science, and network theory, primarily to minimize expected task-completion times modeled via first-passage times (see, e.g., Luby et al. 1993, Avrachenkov et al. 2018, and Evans et al. 2020). In contrast, our work focuses on the asymptotic behavior of the full distribution of passage times, studying large deviation regimes and rare-event phenomena through asymptotic estimates, bounds, and threshold effects rather than expectations. Although related analytic frameworks for Markov processes with restart—addressing ergodicity, quasi-stationary distributions (QSDs), and connections to Fleming-Viot systems—are developed in Grigorescu and Kang (2013), our results concern finer asymptotics for more restricted process classes and different probabilistic questions. Our work also differs from Monthus (2021), which analyzes large deviations for the long-time asymptotics with stochastic restart to a fixed state. In contrast, we consider large deviations with respect to a spatial parameter modeling the complexity of exploring the state space.

Despite the promise of parallel simulation techniques and restart mechanisms, there is a lack of quantitative evidence of their potential benefits. Specifically, the extent to which these strategies improve exploration efficiency and yield better learning outcomes remains unclear. We aim to fill this gap by providing a detailed analysis of the impact of parallel simulations combined with restart mechanisms on the exploration process, modeled as the task of reaching a specific subset of the state space. To achieve this, we use a simplified model based on toy dynamics, specifically random walks and Lévy processes. These processes are chosen for their mathematical tractability and their ability to capture fundamental aspects of the problem, allowing explicit analytical results. Although these dynamics are simple, they provide crucial insights into the mechanisms that govern exploration, forming the basis for future work with more involved Markovian dynamics. By employing this framework, we examine two key aspects of exploration:

- **Exploration complexity**, encapsulated by a parameter that quantifies the fluctuations necessary to achieve effective exploration and reach the target subset of the state space;
- **Exploration diversity**, represented by the number of parallel simulations employed, which increases the breadth of the state-space search and enhances the likelihood of encountering rare, high-reward states within the target subset.

We focus our analysis on a specific regime where the exploration complexity scales linearly with the simulation budget. The toy dynamics framework allows us to examine these parameters within finite time budgets systematically and to quantify how parallel simulations improve efficiency by leveraging diversity, whereas restarting mechanisms enable particles to avoid stagnation by redirecting effort toward more promising regions of the state space, ultimately facilitating the achievement of the exploration goal.

Our approach relies on the systematic analysis of these strategies under finite time budgets, using analytical techniques based on exponential martingales to quantify the effects of parallel simulations and restarting mechanisms. By addressing the interplay between exploration diversity and the power of parallelization and restarting mechanisms, our work provides novel insights into efficient exploration strategies and lays the groundwork for further investigations into more complex stochastic dynamics.

Through this investigation, our goal is not only to demonstrate the potential utility of these techniques but also to delineate their limitations. Understanding these nuances is crucial for developing more efficient stochastic simulations in model-free settings such as RL, particularly in environments characterized by sparse rewards. Ultimately, our findings aim to provide actionable insights into when and how to leverage parallel simulations and restarting strategies to optimize exploration, thereby advancing the field in both theory and practice.

1.1. Problem Statement and Contributions

In this work, we model exploration as a one-dimensional stochastic process, often referred to as a particle, aiming to reach a target subset of the state space. This target, represented as a barrier, captures a potentially informative state that is difficult to attain, making it essential to devise strategies that maximize the probability of success within a finite time budget and to quantify the resulting improvements. In the following we summarize our main findings and key take-away messages.

1. Phase Transition in Reaching Probabilities. We identify a striking phase transition in the probability of reaching the target subset of the state space as a function of the number of parallel simulations N . In Theorems 1 and 2, we rigorously show that when the time budget scales linearly with the barrier level, there is a sharp transition in performance. Our analysis is based on large deviations techniques, and the transition happens at a certain threshold which we present explicitly. Below this threshold, concentrating the entire budget on a single particle is optimal. If the budget, when split into N parts, remains above the threshold, then distributing it among N particles significantly improves the probability of success. However, using too many particles—thereby reducing the time available to each beyond the threshold—results in exponentially diminished performance.

2. Optimal Number of Parallel Simulations. In view of the previous point, a critical contribution is then the determination of an optimal number of parallel simulations, N^* , that balances the trade-off between exploration diversity and the time allocated to each particle. This optimal N^* depends intricately on the time budget and the large deviations characteristics of the particle dynamics, but in terms of the phase transition mentioned above, it can be characterized as the largest integer for which the split budget exceeds the threshold.

3. Exponential Boost via Restarting Strategy. To overcome stagnation and improve the probability of reaching the target, we introduce a restarting mechanism inspired by mechanisms described, for instance, in Mastropietro et al. (2025). We study an idealized situation in which this mechanism replaces particles with a low likelihood of success and restarts them at new initial positions sampled from a given measure. Our theoretical results (Theorem 3) demonstrate that this restarting strategy can substantially enhance the probability of success. The improvement is quantified by a factor proportional to the time budget and the exponential moments of the measure associated with the restarting mechanism. A particularly interesting restarting measure is a QSD over a set representing high-probability regions for successful trajectories. For these restarting measures, we can obtain sharper asymptotics using different methods that exploit the QSD properties. We note that this restarting strategy can be well approximated in practice by selection mechanisms such as the Fleming-Viot or Aldous-Flanary-Palacios algorithms (see Budhiraja et al. 2022 and the references therein). These schemes enable practical implementation and adequate sampling of initial positions, making the restarting mechanism computationally feasible in real-world scenarios. Finally, from a practical point of view, we present contributions in two directions for queueing systems. The first one corresponds to the problem of estimating small probabilities. We analyze the interplay between the probability of exploration and the efficiency of the associated estimation method. For the particular case of an $M/M/1$ queue, we show that the variance of the vanilla Monte Carlo estimator is essentially determined by the typical exploration time. The second, with broader impact, addresses efficient exploration in RL. In RL environments with sparse rewards, the core of the efficient exploration challenge lies in the inability of algorithms such as policy gradient methods to perform effective policy updates (see Ecoffet et al. 2021 and references therein). These methods estimate the gradient of the expected return with respect to the policy parameters. This estimation relies crucially on exploring the state space; if the agent fails to reach rewarding states, the gradient estimates are biased and ineffective. Our work addresses this specific issue: we aim to improve gradient estimation (hence policy evaluation under a fixed policy) not by changing the policy, but by enhancing the exploration process used to evaluate it. The $M/M/1/K$ queue considered in Mastropietro et al. (2025) is reanalyzed here to illustrate our results.

1.2. Organization of the Paper

In Section 2, we introduce the models for the exploration process, which represent the dynamics of the exploration phase in RL. Like the rest of the paper, this section is divided into two parts. The first part analyzes the effects of parallel exploration, and the second part analyzes our restarting mechanism. Section 3 presents our main results for the two scenarios described above, after some necessary preliminaries. Main results are discussed in generality and illustrated through numerical simulations for concrete instances. In Section 4, we prove Theorems 1 and 2

related to parallel exploration. In Section 5 we prove Theorem 3 for exploration with a general restarting measure, and in Section 6 we prove Theorem 4 for a QSD restarting measure. Finally, appendices are devoted to practical applications of our work.

2. Exploration Strategies

We now formally introduce the models for exploration, using random walks and Lévy processes as simplified toy models. We focus on the probability of reaching a distant state x from an initial position of zero, with the drift component encoding the difficulty of exploration. The drift represents the inherent challenge in progressing toward the target state, where a stronger drift away from the target reflects a more difficult exploration task.

We consider scenarios where the total time budget scales with x , allowing us to examine how exploration strategies perform as the difficulty of reaching the target increases. Two specific strategies are explored: the first investigates parallel exploration under a fixed time budget, whereas the second examines the effects of restarting the process when it crosses a predefined threshold. Although these models are simplified, they provide critical insights into the fundamental principles that govern the effectiveness of exploration strategies in RL, particularly in challenging environments where reaching an unlikely target state is the goal.

2.1. Strategy I: Parallel Exploration

Let $\{Z_i(t)\}_{t \in I}, i \in \{1, \dots, N\}$, where $I = \mathbb{N}$ or $I = \mathbb{R}^+$, be a family of one-dimensional random walks or Lévy processes. Define further

$$\tau_i(x) := \inf\{t \in I : Z_i(t) \geq x\}, \quad (1)$$

$$\tau^{(N)}(x) := \min\{\tau_1(x), \tau_2(x), \dots, \tau_N(x)\}, \quad (2)$$

the time it takes particle i to reach state x , and the first time one of N particles reaches x , respectively.

First, the case of N parallel independent simulations is analyzed. We introduce as performance criterion for exploration the probability that starting from zero, one of the simulations reaches the rare state x within a given time budget $B(x)/N$, that is,

$$\mathbb{P}(\tau^{(N)}(x) \leq B(x)/N).$$

We aim at characterizing this probability (as a function of N) in the regime where the total simulation budget, $B(x)$, satisfies $x = C \cdot B(x) + o(\sqrt{B(x)})$, for x large and C some positive constant. This is the focus of Theorem 1 for random walks and Theorem 2 for Lévy processes.

2.2. Strategy II: Exploration with Restart

Our second model incorporates a *restart mechanism* into the dynamics of a random walk or a Lévy process. The ingredients of the model are a positive number $x > 0$ (the rare state we want the process to reach), a probability measure ν_x supported on $(0, x)$, and a discrete random walk or a Lévy process $Z = \{Z(t)\}_{t \in \mathbb{R}^+}$. The trajectories of the process under study can be viewed as a concatenation in time of independent and identically distributed cycles, each consisting of sampling a state y_0 from ν_x and simulating the trajectory of $y_0 + Z(t)$ until the first time it exits the interval $(0, x)$.

Our goal is to analyze the first passage time over x of the restarted process, and compare it with the corresponding one for the process without restart. The conditions required on the family of restart measures to improve algorithmic performance are quite general. Essentially, it suffices that each ν_x be stochastically dominated by a nondegenerate measure possessing a finite second moment. We also examine in detail the case where restart occurs according to the unique quasi-stationary measure of the process Z absorbed at $\mathbb{R}^+ \setminus (0, x)$ (see Subsection 3.4 for precise definitions). Our motivation for this particular mechanism (restarting from quasi-stationary distributions) is that the resulting process approximates complex particle dynamics built upon a given Markov process, a key example, as noted earlier, being the class of Fleming-Viot particle systems. Our model is indeed an approximation of the Fleming-Viot dynamics because the stationary distribution of the N – particle system approximates the quasi-stationary distribution associated with the underlying Markov process absorbed upon reaching A , as N goes to infinity (see Appendix B for more details).

3. Main Results

After introducing due notation, we present our results for each strategy.

3.1. Parallel Exploration Driven by a Random Walk

Let $Z(t) = X_1 + \dots + X_t$, $t \in \mathbb{N}$ be a random walk with independent increments all distributed as a random variable X . We will always assume that X is not deterministic, meaning that it is not concentrated on a single value, and that it may take both positive or negative values with positive probability. The moment-generating function (mgf) of X , $\varphi(\lambda) := \mathbb{E}[\exp(\lambda X)]$, is strictly convex on the interior points of the set Λ ($\text{int}(\Lambda)$) where $\Lambda := \{\lambda \in \mathbb{R} : \varphi(\lambda) < +\infty\}$. The following condition is usually referred to as the *right Cramér condition*:

$$\Lambda_+ := (0, +\infty) \cap \text{int}(\Lambda) \text{ is nonempty.} \quad (3)$$

We will work under further additional assumptions:

$$\varphi \text{ is at least twice differentiable in } \Lambda_+, \quad (4)$$

and

$$\text{There exists } \lambda^* \in \Lambda_+ : \psi(\lambda^*) = 0 \quad \text{where} \quad \psi(\lambda) := \log \varphi(\lambda). \quad (5)$$

Note that $\mathbb{E}[X] = \psi'(0)$ and Assumptions (3), (4), and (5) imply that $\mathbb{E}[X] < 0$ and $\psi'(\lambda^*) > 0$, because $\psi(\lambda^*) = \psi(0)$ and φ is strictly convex. Let $\tau(x)$ denote the first time the random walk reaches level x , defined as

$$\tau(x) := \inf\{t \geq 0 : Z(t) \geq x\}. \quad (6)$$

Furthermore, let $\tau^{(N)}(x)$ represent the minimum of N independent random times, each distributed as $\tau(x)$, that is,

$$\tau^{(N)}(x) := \min\{\tau_1(x), \tau_2(x), \dots, \tau_N(x)\}, \quad (7)$$

which corresponds to the first time any of N independent random walks reaches level x .

Our first main theorem states that the interesting behavior of $\tau(x)$ and $\tau^{(N)}(x)$ occurs while comparing them to functions of the following classes: for $\lambda \in \Lambda_+$ let

$$L(\lambda) := \{f : \mathbb{R} \rightarrow \mathbb{R} \text{ such that } x = \psi'(\lambda) \cdot f(x) + \mathbf{o}(\sqrt{f(x)}) \text{ as } x \rightarrow +\infty\}. \quad (8)$$

This set consists of functions $f(x)$ that asymptotically behave like $x/\psi'(\lambda)$ with an error term of order smaller than $\sqrt{f(x)}$. The result below is stated for the class of functions defined above, but the reader can think of a specific example being $B(x) = x/\psi'(\lambda)$ for any $\lambda \in \Lambda_+$.

We now state our main result regarding parallel exploration for random walks:

Theorem 1. *Let Z denote a random walk whose increments satisfy Assumptions (3), (4), and (5), and let $B(\cdot)$ belong to $L(\lambda)$ for some $\lambda \in \Lambda_+$. Then for $N \geq 2$,*

$$\lim_{x \rightarrow +\infty} \frac{\mathbb{P}\left(\tau^{(N)}(x) \leq \frac{B(x)}{N}\right)}{\mathbb{P}\left(\tau(x) \leq B(x)\right)} = \begin{cases} N, & \text{if } N\psi'(\lambda) < \psi'(\lambda^*); \\ 0, & \text{if } N\psi'(\lambda) > \psi'(\lambda^*). \end{cases} \quad (9)$$

Corollary 1 (Optimal Number of Particles). *For a given state x and a budget $B(x) \in L(\lambda)$. If $\psi'(\lambda) < \psi'(\lambda^*)$, the asymptotically optimal number of particles is*

$$N^* = \max\{N \geq 1 : N\psi'(\lambda) < \psi'(\lambda^*)\} = \left\lceil \frac{\psi'(\lambda^*)}{\psi'(\lambda)} \right\rceil - 1, \quad (10)$$

where $\lceil \cdot \rceil$ is the ceiling function. Otherwise, $N^* = 1$.

In the regime in which the theorem is stated, probabilities decay exponentially as x goes to infinity (see Proposition 1 in Section 4). The result shows that using too many particles (thus giving too little time to each) entails an exponentially worse performance, meaning that $\mathbb{P}(\tau^{(N)}(x) \leq \frac{B(x)}{N})$ decays with a faster exponential rate than $\mathbb{P}(\tau(x) \leq B(x))$. On the other hand, using too few particles comes only at a linear cost in performance.

The main intuition behind the results stated above is that within the (rare) event $\{\tau(x) < +\infty\}$, the probability is concentrated on those trajectories with drift $\psi'(\lambda^*)$, for which the time needed to reach x is $\tau(x) \approx x/\psi'(\lambda^*)$. This idea is formalized using exponential martingales and is at the core of Proposition 1. Our results show that allowing a budget smaller than $x/\psi'(\lambda^*)$ substantially reduces the probability of success, whereas allowing too much budget results in no significant improvement. In the latter case, by splitting it into several independent particles, each one provides an “extra chance” of success.

Even though in general λ^* and (thus) the threshold $\psi'(\lambda^*)^{-1}$ depend in a nontrivial way on the distribution of the increments, in some cases it is possible to find simple expressions for these parameters. We present some of the important examples here.

Example 1 (Normal Increments). If X has a normal distribution with mean $-\mu < 0$ and variance $\sigma^2 = 1$, then

$$\psi(\lambda) = -\lambda\mu + \frac{1}{2}\lambda^2;$$

hence, $\lambda^* = 2\mu$ and $\psi'(\lambda^*) = \mu$ and the threshold can be reliably estimated in practice if needed. Moreover, the function $\psi' : \Lambda_+ \rightarrow [0, +\infty)$ is surjective, so for every $C > 0$ there exists $\lambda \in \Lambda_+$ such that $\psi'(\lambda) = 1/C$. Theorem 1 in this case states that if, in order to reach a high state $x > 0$, we are allowed a budget $B(x) = Cx$ for some $C > 0$, then the (asymptotically) optimal number of independent particles to use is $N^* = \lceil C\mu \rceil - 1$ if $C\mu > 1$ and $N^* = 1$ if not.

Example 2 (Birth-and-Death Chains). If X only takes values 1 and -1 with probabilities p and $1 - p$, respectively, with $p < 1/2$, then $\lambda^* = \log(1 - p) - \log p$ and $\psi'(\lambda^*) = 1 - 2p = -\mathbb{E}X$. Because X is bounded, the image of ψ' contains the positive real numbers as in the previous example. Then, if we are allowed a budget $B(x) = Cx$, the number of particles maximizing our chances is $N^* = \lceil C(1 - 2p) \rceil - 1$ if $C(1 - 2p) > 1$ and $N^* = 1$ if not.

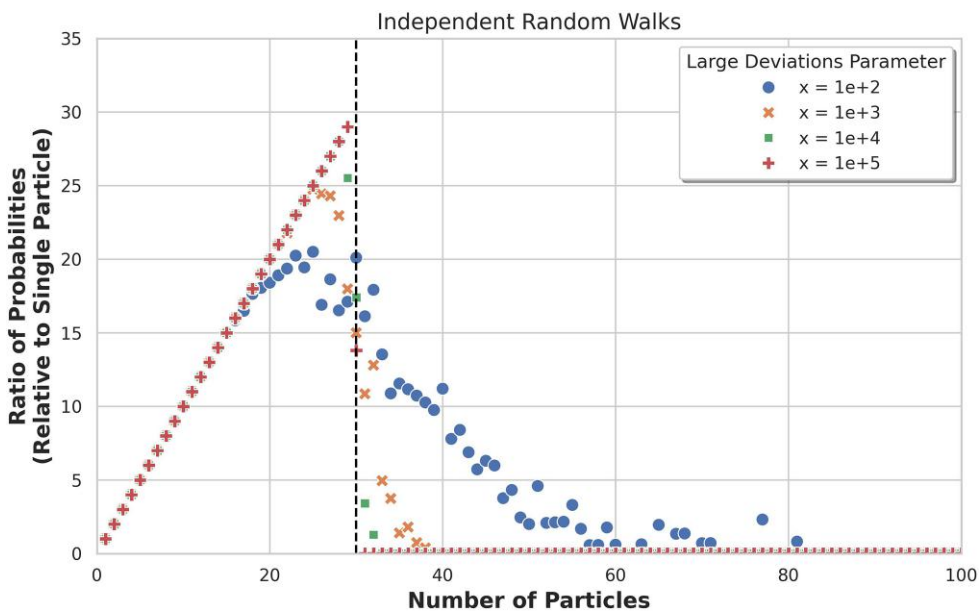
3.1.1. Simulations for Parallel Random Walks. Consider the random walk case described in Example 2. We use the parameters $p = 0.45$ and $B(x) = 300 \cdot x$, which imply $\psi'(\lambda^*) = 0.1$. According to Theorem 1 the ratio

$$\frac{\mathbb{P}\left(\tau^{(N)}(x) \leq \frac{B(x)}{N}\right)}{\mathbb{P}\left(\tau(x) \leq B(x)\right)} \quad (11)$$

is expected to approach the identity for $N < 300 \cdot 0.1 = 30$ and zero for $N > 30$, as $x \rightarrow \infty$. This is the behavior observed in Figure 1.

For the simulation we used the measure \mathbb{P}^{λ^*} as defined in (20), under which the process is distributed as a birth-and-death random walk with interchanged parameters, namely, $p^* = \mathbb{P}^{\lambda^*}(X_1 = 1) = 1 - \mathbb{P}^{\lambda^*}(X_1 = -1) = 1 - p$. Under this measure, passage over a given $x \in \mathbb{N}$ is not a rare event anymore, so simulations are feasible. For each large deviation parameter $x = 10^k$, $k \in \{2, 3, 4, 5\}$, and for each number of particles N between 1 and 100, we simulate and average over 1,000 random walks to get an estimate of $\mathbb{P}^{\lambda^*}(\tau \leq B(x)/N)$. By reverting the measure change, we obtain estimates for the original ratio.

Figure 1. Parallel Exploration for Random Walks with Negative Mean $p - (1 - p) = -0.1$



Notes. Estimated ratio between $\mathbb{P}(\tau^{(N)}(x) \leq B(x)/N)$ and $\mathbb{P}(\tau(x) \leq B(x))$ as a function of number N with $B(x) = C \cdot x = 300 \cdot x$. The phase transition is observed at the expected threshold $N^* = \lceil C(1 - 2p) \rceil - 1 = 29$.

3.2. Parallel Exploration Driven by a Lévy Process

Random walks have an analogue in continuous time known as Lévy processes. Under analogous assumptions on the exponential moments, we can extend Theorem 1 to the case of Lévy processes on \mathbb{R} . This is based on the fact that the asymptotics for first passage times over high barriers are similar in the cases of random walks and Lévy processes; see Section 4. We now review some preliminaries on Lévy processes, all of which may be found in classical references such as Kyprianou (2006) and Bertoin (1998).

Let $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_t, \mathbb{P})$ be a filtered probability space, where the filtration $\{\mathcal{F}_t\}_t$ is assumed augmented and right-continuous. A Lévy process $Z = \{Z(t)\}_{t \in I}$, $I = [0, +\infty)$ is a stochastic process such that: for all $0 \leq s \leq t$, $Z(t)$ is \mathcal{F}_t measurable, $Z(0) = 0$, $Z(t) - Z(s)$ is independent of \mathcal{F}_s and has the same distribution as $Z(t - s)$, and is continuous in probability, namely, that $Z(t + s) \rightarrow Z(t)$ in probability if $s \rightarrow 0$. We will always work with a càdlàg version of Z . Let ψ denote the Lévy exponent of Z :

$$\mathbb{E}e^{\lambda Z(t)} = e^{t\psi(\lambda)}.$$

The Lévy-Khintchine formula introduces the characteristic triplet $(-\mu, \sigma, \Pi)$ of Z :

$$\psi(\lambda) = \log \mathbb{E}e^{\lambda Z(1)} = -\mu\lambda + \frac{(\sigma\lambda)^2}{2} + \int_{-\infty}^{+\infty} (\exp(\lambda y) - 1 - \lambda y \mathbf{1}(|y| \leq 1)) \Pi(dy), \quad (12)$$

where $\mu \in \mathbb{R}$ is called the drift coefficient, $\sigma \in (0, +\infty)$ is the diffusion coefficient (note that we exclude the finite-variation case), and Π is a measure in $\mathbb{R} \setminus \{0\}$ such that

$$\int_{-\infty}^{+\infty} (1 \wedge y^2) \Pi(dy) < +\infty,$$

called the Lévy measure (or jump measure). As in the discrete-time case we will work under Cramér's condition: $\mathbb{E}(e^{\lambda Z(1)})$ is finite for $\lambda \in [0, \lambda_{\max})$ and there exists $\lambda^* \in (0, \lambda_{\max})$ such that $\psi(\lambda^*) = 0$. This condition implies in particular that $\mathbb{E}Z(1) < 0$ and that Z drifts to $-\infty$ almost surely.

Let $\tau(x)$ and $\tau^{(N)}(x)$ be defined as (6) and (7), respectively. Up to taking a continuous time parameter, our result on the number of particles under finite time budget constraints for Lévy processes is the same as for random walks:

Theorem 2. Let Z denote a Lévy process such that $Z(1)$ satisfies Assumptions (3), (4), and (5). Let $B(\cdot)$ belong to the class $L(\lambda)$ for some $\lambda \in \Lambda_+$ (as defined in (8)). Then for $N \geq 2$,

$$\lim_{x \rightarrow +\infty} \frac{\mathbb{P}\left(\tau^{(N)}(x) \leq \frac{B(x)}{N}\right)}{\mathbb{P}\left(\tau(x) \leq B(x)\right)} = \begin{cases} N, & \text{if } N\psi'(\lambda) < \psi'(\lambda^*); \\ 0, & \text{if } N\psi'(\lambda) > \psi'(\lambda^*). \end{cases} \quad (13)$$

In view of the above result, we obtain an optimal number of particles for the Lévy process, in the same way as in Corollary 1.

The following is an example where all the parameters can be computed explicitly.

Example 3 (Linear Brownian Motion). If the Lévy measure is zero everywhere, the process is a Brownian motion with drift, namely, a solution of the stochastic differential equation (SDE)

$$dZ(t) = -\mu dt + \sigma dB(t).$$

We assume further that $\sigma = 1$. Because its distribution at time 1 is Gaussian with parameters $(-\mu, 1)$, from Example 1 we conclude that $\lambda^* = 2\mu$ and $\psi'(\lambda^*) = \mu$. The same conclusion on the optimal number of particles as in Example 1 holds.

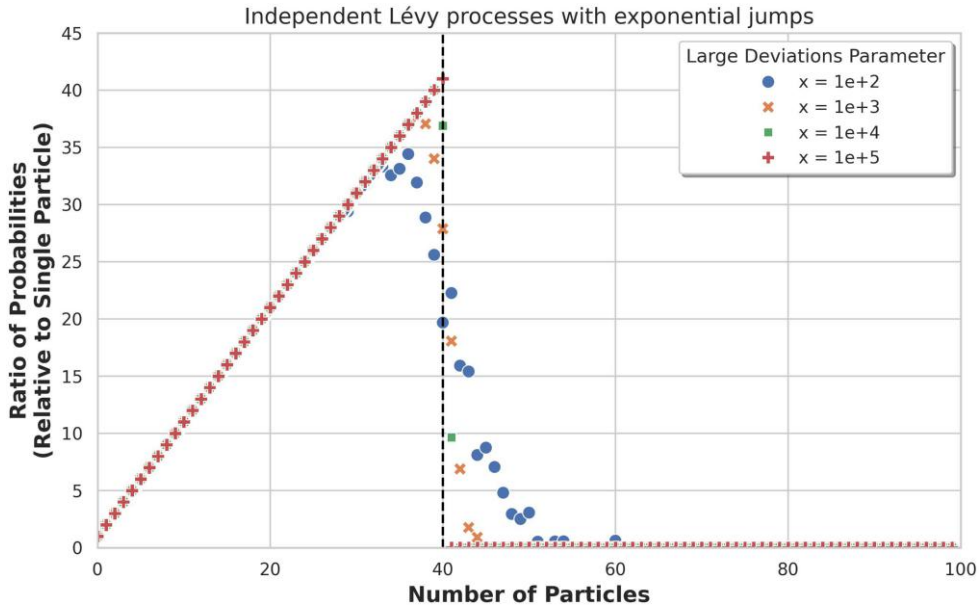
3.2.1. Simulations for Parallel Lévy Processes. For the simulations in Figure 2 we considered a Lévy process with positive and negative jumps. Positive (resp. negative) jumps occur at times distributed as a Poisson process of intensity $r = 2$ (resp. $s = 3$) and whose lengths are exponentially distributed with rate $\alpha = 4$ (resp. $\beta = 1$). Following the notation in (12), we further take $\mu = \sigma = 1$. It can be checked that the Lévy measure is given by

$$\Pi(dy) = \mathbf{1}(y < 0)3 \exp(y) dy + \mathbf{1}(y > 0)8 \exp(-4y) dy.$$

For the considered process, $\psi(\lambda) = -\lambda + \frac{1}{2}\lambda^2 - \frac{3\lambda}{\lambda+1} + \frac{2\lambda}{4-\lambda}$, $\lambda < 4$ and its positive root turns out to be $\lambda^* = 2$. Moreover, $\psi'(2) = 8/3$, which gives, for B taken as $B(x) = 15 \cdot x$, an optimal number $N^* = 40$ of particles.

For simulation purposes, as in the discrete-time case, we use the exponential martingale associated with λ^* and the measure \mathbb{P}^{λ^*} , as defined in Section 4. Specifically, the parameters of the process were chosen in such a way that

Figure 2. Parallel Exploration Lévy Processes



Note. Estimated ratio between $\mathbb{P}(\tau^{(N)}(x) \leq B(x)/N)$ and $\mathbb{P}(\tau(x) \leq B(x))$ as a function of the number N of independent Lévy processes with exponential jumps, with parameters described in Subsection 3.2.1.

under the measure \mathbb{P}^{λ^*} , the process is again Lévy with exponential jumps. Moreover, the time intensities and the jump rates are permuted: under the measure \mathbb{P}^{λ^*} the parameters of the Lévy measure are $r^* = 4$, $\alpha^* = 2$, $s^* = 1$, and $\beta^* = 3$.

Remark 1 (A Remark on Reflected Process). In terms of applications, it may be desirable to include reflection at a low barrier in our models to convey the idea that complexity of an exploration task cannot diverge to $-\infty$ in realistic scenarios. However, in the context of rare events we are in, there is no substantial difference in introducing reflection, because estimates vary at most by a multiplicative factor. To see this, consider the process $\tilde{Z}(t) = Z(t) - [\inf_{0 \leq s \leq t} Z(s) \wedge (-1)]$, namely, the reflected version of Z at -1 . By considering the coupling $(Z(t), \tilde{Z}(t))$ and observing that $\mathbb{P}(Z(t) \leq \tilde{Z}(t)) = 1$ for every $t \geq 0$, it follows that

$$\mathbb{P}(\tau(x) \leq t) \leq \mathbb{P}(\tilde{\tau}(x) \leq t), \text{ for every } t \geq 0.$$

On the other hand, the event $\{\tilde{\tau}(x) < t\}$ may be decomposed into trajectories that reach x before the first reflection at -1 and those that do not. The first subset may be treated as a process without reflection, thus contributing a term of the same order as that of the process without reflection. Trajectories of the latter subset have a last reflection time before $\tilde{\tau}(x)$, so there is an excursion of length $x + 1$ during a shorter time. This provides an extra term of order at most $\mathbb{P}(\tau(x + 1) \leq t) < \mathbb{P}(\tau(x) \leq t)$. Thus,

$$\mathbb{P}(\tilde{\tau}(x) \leq t) \leq 2\mathbb{P}(\tau(x) \leq t),$$

and we may then safely restrict ourselves to processes without reflection. The interested reader may consult Hansen (2009) for the maximum of reflected Lévy processes.

3.3. Exploration with Restart for Random Walks and Lévy Processes

Our Strategy II consists of restarting a random walk or a Lévy process such as those studied in Theorems 1 and 2 upon leaving $(0, x)$ with a given probability measure ν_x supported in the interval $(0, x)$. Our first result in this direction, stated in Theorem 3, gives asymptotic performance guarantees of restart strategies under mild assumptions on the restart measures ν_x , as $x \rightarrow +\infty$.

In the next section we will specialize to a particularly interesting family of measures: quasi-stationary measures, for which better estimates are obtained (in particular matching bounds up to multiplicative constants). We remark that for the QSD case we rely exclusively on quasi-stationarity and related properties but not on Theorem 3.

Definition 1. Let ν_1, ν_2 be two positive finite measures on \mathbb{R} . ν_1 is said to be *stochastically dominated* by ν_2 , denoted $\nu_1 \preceq_{\text{stoch}} \nu_2$, if for every measurable nondecreasing function u ,

$$\int_{\mathbb{R}} u(y) \nu_1(dy) \leq \int_{\mathbb{R}} u(y) \nu_2(dy). \quad (14)$$

This notion is sometimes called *first-order stochastic domination* and is equivalent to $\nu_1(x, +\infty)$ being less than or equal to $\nu_2(x, +\infty)$ for every $x \in \mathbb{R}$.

The next theorem provides estimates on the time it takes for the restarted process to reach the desired region (here an interval $[x, +\infty)$ for large x) under mild assumptions on the restart measures $(\nu_x)_{x \geq 1}$, for which we assume without loss of generality that $x \geq 1$.

Theorem 3. Let $\{Z(t)\}_{t \geq 0}$ denote a random walk or a Lévy process satisfying Cramér’s condition, and let λ^* be given by (5). For each $x > 0$, let $\{Z^{\nu_x}(t)\}_{t \geq 0}$ be the restarted process with restart measures $(\nu_x)_{x \geq 1}$ as described at the beginning of this section, and let $\tau^{\nu_x}(x)$ be the first time a cycle of the process ends above x . Assume that there exists a finite positive measure ν with the following properties:

1. ν is not a delta measure on zero;
2. $\nu_x \preceq_{\text{stoch}} \nu$ for every positive x ;
3. ν has a finite second moment: $\int y^2 \nu(dy) < +\infty$.

Then, $B(x)$ growing faster than $\psi'(\lambda^*)^{-1}x$ but slower than e^{λ^*x} , the ratio

$$\frac{\mathbb{P}(\tau^{\nu_x}(x) \leq B(x))}{\mathbb{P}(\tau(x) \leq B(x))} \frac{1}{B(x) \int_0^x \exp(\lambda^*y) \nu_x(dy)} \quad (15)$$

is bounded away from zero uniformly in $x > 0$.

A typical trajectory of the restarted process is a concatenation of independent cycles, each of which consists of a trajectory $\{y_0 + Z(t) : 0 \leq t \leq \tau\}$ for some $y_0 \in (0, x)$ sampled according to ν_x and $\tau := \inf\{t > 0 : y_0 + Z(t) \notin (0, x)\}$. An intuition behind the result is that adding ν_x as an initial distribution boosts the performance by a factor $\int_0^x \exp(\lambda^*y) \nu_x(dy)$ for each cycle. Because the cycles’ duration is stochastically bounded, we expect to observe an asymptotically linear number of cycles within time $B(x)$. However, a matching upper bound remains elusive for this general context.

Example 4. Consider some fixed probability measure ν on $(0, +\infty)$ such that $\nu((0, 1)) > 0$ and satisfying properties 1, 2, and 3 listed in Theorem 3. Now define the sequence $(\nu_x)_{x \geq 1}$ as the conditional laws

$$\nu_x(\cdot) := \frac{\nu(\cdot \cap (0, x))}{\nu((0, x))}.$$

Then $\nu_x(y, +\infty) \leq \nu((0, 1))^{-1} \nu(y, +\infty)$, and Theorem 3 is applicable for this family of measures.

Remark 2. In a more general fashion than the previous example, one may wonder if the hypotheses of Theorem 3 are satisfied whenever the measures ν_x have a weak limit ν . This is indeed the case under some assumptions on the moments of the ν_x ’s, but the dominating measure need not be related to the weak limit. Define $\nu_*((-\infty, y]) := \inf_{x \geq 1} \nu_x((-\infty, y])$. Then ν_* defines a càdlàg function with $\lim_{y \rightarrow -\infty} \nu_*((-\infty, y]) = 0$. Because the ν_x ’s have a weak limit, $\lim_{y \rightarrow \infty} \nu_*((-\infty, y]) = 1$ also holds and ν_* defines a probability measure that dominates ν_x for every $x \geq 1$. Note as well that ν_* is not a delta measure at zero unless $\nu_x = \delta_0$ for all $x \geq 1$. Let us now give a precise condition under which ν_* satisfies property 3 in Theorem 3. Assume that the ν_x ’s have uniformly bounded moments of order p for some $p > 2$, that is,

$$C := \sup_{x \geq 1} \int_0^{+\infty} y^p \nu_x(dy) < +\infty. \quad (16)$$

We bound the tails of the ν_x by Markov’s inequality:

$$\nu_x(y_0, +\infty) \leq y_0^{-p} \int_0^{+\infty} y^p \nu_x(dy) \leq C y_0^{-p}.$$

Then, using the Layer Cake Representation, the desired property holds:

$$\begin{aligned} \int_0^{+\infty} y^2 \nu_*(dy) &= 2 \int_0^{+\infty} y \nu_*(y, +\infty) dy = 2 \int_0^{+\infty} y \sup_{x \geq 1} \nu_x(y, +\infty) dy \\ &\leq 2C \int_0^{+\infty} y^{1-p} dy < +\infty. \end{aligned}$$

As a final example for this section, we take a simple extension of Example 4 that shows that Theorem 3 covers cases beyond weak convergence:

Example 5. Take ν^1 and ν^2 two (distinct) probability measures on $(0, +\infty)$ satisfying the conditions of Theorem 3 and $\nu^i((0, 1)) > 0$, $i = 1, 2$. We may build an alternating family of measures for $x \geq 1$ as follows:

$$\nu_x(\cdot) := \begin{cases} \frac{\nu^1(\cdot \cap (0, x))}{\nu^1((0, x))} & \text{if } \lfloor x \rfloor \text{ is odd,} \\ \frac{\nu^2(\cdot \cap (0, x))}{\nu^2((0, x))} & \text{if } \lfloor x \rfloor \text{ is even.} \end{cases}$$

In this case Theorem 3 applies, but there is not a weak limit of the whole sequence of measures.

3.4. Restart with Quasi-stationary Measures

We now specialize to an important family of restarting measures for which we can obtain sharper asymptotic results than those provided by Theorem 3.

Throughout this section we assume that ν_x is the QSD of the process restricted to the interval $(0, x)$. We begin with a few brief preliminaries on quasi-stationary distributions, and refer the reader to the monograph of Collet et al. (2013) for a comprehensive treatment of the subject.

Definition 2 (Quasi-stationary Measure). Let $Z = \{Z(t)\}_{t \geq 0}$ be a Markov process with state space \mathcal{S} and let A be a subset of \mathcal{S} . The absorbed process, denoted by Z_A , is defined as $Z_A(t) = Z(t \wedge \tau(A))$, where $\tau(A) := \inf\{t > 0 : Z(t) \in A\}$. A quasi-stationary measure (QSD) ν is a probability measure on $\mathcal{S} \setminus A$ which is invariant when conditioned on nonabsorption, which means

$$\mathbb{P}_\nu(Z(t) \in B \mid t < \tau(A)) = \frac{\mathbb{P}_\nu(Z_A(t) \in B)}{\mathbb{P}_\nu(t < \tau(A))} = \nu(B),$$

for every $t > 0$ and every measurable set $B \subset \mathcal{S} \setminus A$.

The study of QSD for a given Markov process is a subtle matter; in general, they may not exist, and when they do, they may not be unique.

For the particular case of a one-dimensional Lévy process with absorbing set $A = \mathbb{R} \setminus (0, a)$ for some $a > 0$, by Kolb and Savov (2014, theorem 2.1 and remark 2.2) there is a unique QSD, provided the distribution of the process without absorption at any time $t > 0$ is absolutely continuous with respect to Lebesgue in \mathbb{R} , which is satisfied because we are assuming a strictly positive diffusion coefficient. When there is a unique QSD, it is obtained as *the Yaglom limit*

$$\lim_{t \rightarrow +\infty} \mathbb{P}_y(Z(t) \in \cdot \mid t < \tau(A)),$$

which is independent of the starting position $y \in \mathcal{S} \setminus A$.

The existence of a Yaglom limit for a Lévy process when the absorbing set is $(-\infty, 0]$ was proved in Kyprianou and Palmowski (2006),¹ extending the case without jumps of Martinez and San Martin (1994).

To the best of our knowledge, there is no analogous existence-and-uniqueness result specific to discrete-time random walks conditioned on nonabsorption on bounded intervals. If the law of the increments of a random walk is absolutely continuous with respect to the Lebesgue measure, then one may consider the associated compound Poisson process (CPP) with rate 1, thereby obtaining a Lévy process with an absolutely continuous law for every time larger than its first jump. By checking that for large t the probability of not having jumped is negligible for the CPP *even when conditioned on nonabsorption by time t* , one obtains the same Yaglom limit as the conditioned CPP and in particular the existence and uniqueness given by Kolb and Savov (2014). This can be carried out using the asymptotics on passage times from Denisov and Shneer (2013). These asymptotics for conditioned processes are used in the proof of Lemma 4 of Section 6. Yaglom limits for random walks with absorption at $(-\infty, 0]$ were studied in Iglehart (1974), and under the Cramér condition in Doney (1985).

We may now state our main theorem for restarted processes:

Theorem 4. Let Z be a random walk on discrete time or a Lévy process satisfying (3), (4), and (5) as in Theorem 1. For each positive $x \in (0, +\infty)$, let ν_x denote the quasi-stationary probability measure on $(0, x)$ for the process Z absorbed upon exiting $(0, x)$. Let Z^{ν_x} be the restarted version of Z on $(0, x)$ with **restart** measure ν_x and $\tau^{\nu_x}(x)$ the associated passage time over x . Let the time budget $B(x)$ grow faster than $x/\psi'(\lambda^*)$ as $x \rightarrow +\infty$ but slower than e^{λ^*x} . Then there exist constants $c, C > 0$

such that

$$c \leq \frac{\mathbb{P}(\tau^{v_x}(x) < B(x))}{\mathbb{P}(\tau(x) < B(x))} \cdot \frac{1}{B(x) \int_0^x e^{\lambda^* y} v_x(dy)} \leq C \quad (17)$$

for all $x > 0$.

Corollary 2. *In the particular case when Z is a linear Brownian motion, with a negative drift equal to $-\mu$, we have*

$$\int_0^x e^{\lambda^* y} v_x(dy) = e^{\mu x}, \quad (18)$$

and the ratio

$$\frac{\mathbb{P}(\tau^{v_x}(x) < B(x))}{\mathbb{P}(\tau(x) < B(x))} \cdot \frac{1}{B(x)e^{\mu x}} \quad (19)$$

is bounded away from zero and infinity for $B(x)$ growing faster than x/μ .

The theorem states that the application of a restarting mechanism with quasi-stationary distributions improves performance by a factor comparable with the time budget times an exponential moment of the respective measure. In the Brownian motion case, this improvement factor is exponential in the complexity of the task. Indeed, although $\mathbb{P}(\tau(x) \leq B(x))$ is of order $e^{-2\mu x}$, with the restarting mechanism we get order $B(x)e^{-\mu x}$. For a general Lévy process, exponential moments of its quasi-stationary measures are not easy to estimate. We conjecture that they grow as $\exp[(\lambda^* - \lambda_0)x]$, as $x \rightarrow +\infty$, where λ_0 is the positive solution of $\psi'(\lambda) = 0$. Note that this is exactly the case for the linear Brownian motion. For spectrally negative Lévy processes, namely, when $\Pi(0, +\infty) = 0$, the conjecture might be proved with the use of a Girsanov transformation of the process into a zero-drift one (thus providing the conjectured exponential factor) and applying the results of Lambert (2000). We do not pursue this further here; it is left as a direction of further study.

3.4.1. Simulations for Restarted Lévy Processes. We now present numerical experiments for restarted processes. As a case study, we consider the Lévy process with exponential jumps introduced in Subsection 3.2.1, equipped with a restart mechanism activated upon exiting the interval $(0, 50) \subset \mathbb{R}$.

For the process without restart, recall that the Cramér exponent is $\lambda^* = 2$. Consequently, when the process is started from the origin, the probability of ever reaching the upper barrier at 50 is of order $\exp(-100)$. As restart measure, we consider an exponential distribution with mean 10, truncated to $(0, 50)$, namely,

$$v_{50}(dx) = \frac{0.1 e^{-0.1x}}{1 - e^{-5}} \mathbf{1}_{(0,50)}(x) dx.$$

Its exponential moment of order λ^* satisfies

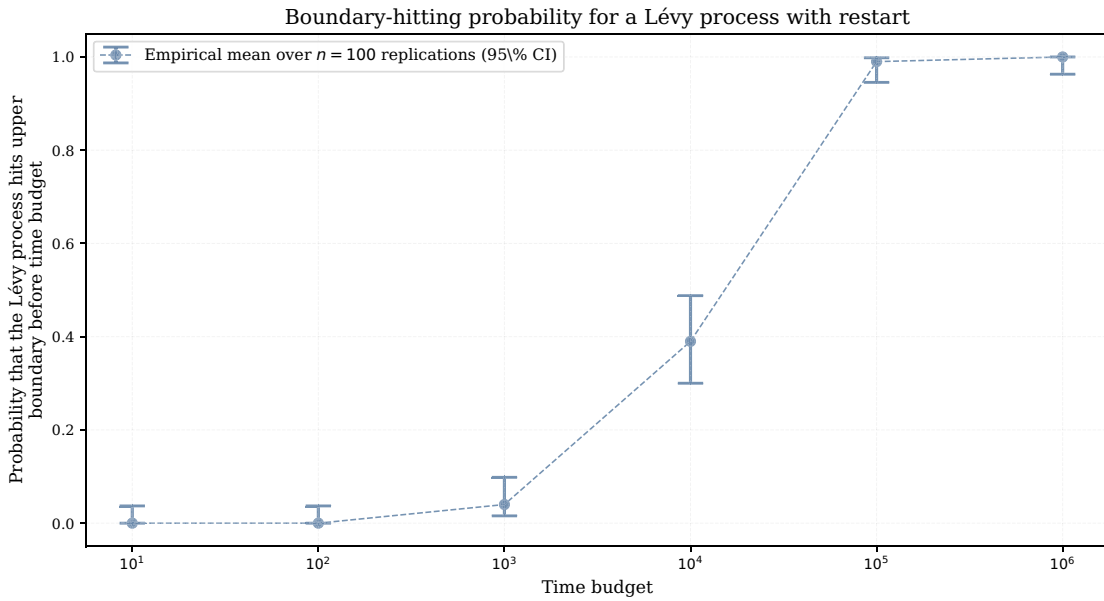
$$\int_0^{50} e^{2x} v_{50}(dx) \approx 9.6 \times 10^{39}.$$

Multiplying by $\mathbb{P}_0(\tau(50) < +\infty)$ yields

$$\mathbb{P}_0(\tau(50) < +\infty) \int_0^{50} e^{2x} v_{50}(dx) \approx 3.6 \times 10^{-4}.$$

Although this estimate is only heuristic, in view of Theorem 3 it suggests that the probability of exceeding the level $x = 50$ within a finite time horizon $B(50)$ under this restart strategy is no longer negligible. This behavior is confirmed numerically in Figure 3, which reports empirical mean estimates based on 100 independent replications for each time budget, together with 95% confidence intervals. The figure illustrates how the restart mechanism makes the exceedance of the upper barrier observable on moderate time scales, with confidence intervals (CI) indicating the associated variability.

We emphasize that the simulations reported in this section do not rely on importance sampling or on any other variance-reduction technique. The numerical results therefore illustrate that introducing a restart mechanism alone can already lead to a significant improvement in the simulation of rare events. In particular, restarting makes it possible to turn probabilities that are extremely small for the original process into quantities that can be reliably estimated within reasonable computational time.

Figure 3. Probability of Exceeding the Level $x = 50$ for the Restarted Lévy Process as a Function of the Time Horizon

Notes. The underlying process combines a Brownian motion with drift $\mu = -1$ and volatility $\sigma = 1$ with exponential jumps: positive jumps arrive at rate $\lambda_+ = 2$ with jump sizes of rate 4, whereas negative jumps arrive at rate $\lambda_- = 3$ with jump sizes of rate 1. Whenever the process exits the interval $(0, 50)$, it is restarted according to a truncated exponential distribution with mean 10.

4. Proof of Theorems 1 and 2

We begin the section by introducing some preliminaries on exponential martingales, first for random walks and then for the Lévy process, which will be used in the proofs. They are standard techniques for the study of large deviations events for light tails, as is our case.

4.1. Preliminaries

By Assumption (3), the family of exponential martingales indexed by $\lambda \in \Lambda_+$ may be defined as

$$M_t^\lambda := \exp(\lambda Z(t) - t\psi(\lambda)), \quad t \in \mathbb{N}.$$

Indeed, M^λ is a martingale for every $\lambda \in \Lambda_+$, which acts as a density for a new probability measure on the same probability space. Specifically, we define the measure \mathbb{P}^λ as

$$\mathbb{P}^\lambda(A) := \mathbb{E}[M_\tau^\lambda \mathbf{1}(A)], \quad (20)$$

where A depends on $\{Z(t)\}_{t \leq \tau}$, and τ is any almost-surely finite stopping time. The measure \mathbb{P}^{λ^*} , associated with λ^* in Equation (5), will be of particular importance.

We also observe that the density of $\mathbb{P}|_{\mathcal{F}_\tau}$ with respect to $\mathbb{P}|_{\mathcal{F}_\tau}$ is given by $(M_\tau^\lambda)^{-1}$.

Exponential martingales are particularly useful for computing probabilities of rare events, as these events may become more frequent under \mathbb{P}^λ if λ is appropriately chosen. Under the measure \mathbb{P}^λ , the increments X_j , $j \geq 1$, are i.i.d. (making Z a random walk), with mean and variance given by

$$\mu(\lambda) := \mathbb{E}_\lambda[X_1] = \psi'(\lambda) \quad \text{and} \quad \sigma(\lambda) := \mathbb{E}_\lambda[|X_1 - \mu(\lambda)|^2] = \psi''(\lambda).$$

The threshold in Theorem 1 can thus be interpreted in terms of the measure \mathbb{P}^{λ^*} : the random time $\tau(x)$ is \mathbb{P}^{λ^*} -almost surely finite (because $\psi'(\lambda^*) > 0$) with mean $x\psi'(\lambda^*)^{-1}$, asymptotically as $x \rightarrow +\infty$.

For the Lévy process we have analogous definitions: under Cramér's condition the family $\{M_t^\lambda\}_{0 \leq t \leq \lambda_{\max}}$ of exponential martingales is defined as

$$M_t^\lambda := \exp(\lambda Z(t) - t\psi(\lambda)),$$

and the corresponding family of measures is given by

$$\left. \frac{d\mathbb{P}^\lambda}{d\mathbb{P}} \right|_{\mathcal{F}_t} = M_t^\lambda.$$

Under any of the measures \mathbb{P}^λ , Z is again a Lévy process with Lévy exponent

$$\psi^{(\lambda)}(r) = \psi(\lambda + r) - \psi(\lambda), \tag{21}$$

with mean and variance again given by

$$\mu(\lambda) := \mathbb{E}_\lambda[X_1] = \psi'(\lambda) \quad \text{and} \quad \sigma(\lambda) := \mathbb{E}_\lambda[|X_1 - \mu(\lambda)|^2] = \psi''(\lambda).$$

From (21) it is straightforward to check that the corresponding characteristic triplet is given by

$$\mu_\lambda = -\mu + \lambda\sigma^2 + \int_{-1}^1 y(e^{\lambda y} - 1)\Pi(dy), \quad \sigma_\lambda^2 = \sigma^2, \quad \Pi^\lambda(dy) = e^{\lambda y}\Pi(dy). \tag{22}$$

The measure \mathbb{P}^{λ^*} provides an intuition to Theorem 2 as in the random walk case: the law of the process conditioned to reach the high barrier is distributed as \mathbb{P}^{λ^*} . We do not prove this exact result here, but it lies at the core of the references cited for our proof (Bertoin and Doney 1994, Palmowski and Pistorius 2009).

Remark 3. μ_λ in (22) should not be confused with $\mu(\lambda)$. If the jump measure Π is nonzero, they might not coincide.

4.2. Proofs

Remark 4. We introduce some pieces of notation: $f(x) \sim g(x)$ will be used to denote asymptotic equivalence $\lim_{x \rightarrow +\infty} f(x)g(x)^{-1} = 1$ and $f(x) \asymp g(x)$ for $C^{-1} < f(x)g(x)^{-1} < C$ for some constant $C > 0$ for all $x > 0$.

Theorems 1 and 2 will follow from the following asymptotics of the passage times of random walks and Lévy processes with the assumptions on their exponential moments stated in Subsection 3.1.

The following proposition unifies Höglund (1990, corollary 2.2) for the case of random walks and Palmowski and Pistorius (2009, theorem 1) for Lévy processes:

Proposition 1. Let $Z = \{Z(t)\}_{t \in I}$ be either a random walk ($I = \mathbb{N}$) or a Lévy process ($I = \mathbb{R}^+$) taking values in \mathbb{R} and assume that Conditions (3), (4), and (5) are satisfied (for $Z(1)$ in the case of a Lévy process). Fix some $\lambda \in (0, +\infty) \cap \Lambda$ such that $\mu(\lambda) = \psi'(\lambda) > 0$. If x and $t = t(x)$ go to infinity in such a way that $t \in L(\lambda)$, then there exist positive constants C (independent of λ) and $D(\lambda)$ such that

$$\mathbb{P}(\tau(x) \leq t) \sim \begin{cases} C \exp(-\lambda^* x), & \text{if } \mu(\lambda) < \mu(\lambda^*) \\ D(\lambda) t^{-1/2} \exp(-\zeta[\mu(\lambda)]t), & \text{if } \mu(\lambda) > \mu(\lambda^*), \end{cases} \tag{23}$$

where ζ is the convex conjugate of ψ :

$$\zeta[s] := \sup_{\lambda \in \Lambda} \{\lambda s - \psi(\lambda)\}. \tag{24}$$

Some remarks are now in order:

Remark 5. The constants C and $D(\lambda)$ are given in Höglund (1990) and Palmowski and Pistorius (2009) in terms of the increments of the process in each case. Their explicit values are not relevant for our analysis.

Remark 6. If $s = \mu(\lambda)$, then λ realizes the supremum in the definition of $\zeta[s]$, that is,

$$\zeta[\mu(\lambda)] = \lambda\mu(\lambda) - \psi(\lambda), \quad \lambda \in \Lambda. \tag{25}$$

Remark 7. The case $\mu(\lambda) < \mu(\lambda^*)$ in the proposition shows that for t growing sufficiently fast with x , $\mathbb{P}(\tau(x) \leq t)$ and $\mathbb{P}(\tau(x) < +\infty)$ have the same exponential order:

$$\begin{aligned} \mathbb{P}(\tau(x) \leq t) &\leq \mathbb{P}(\tau(x) < +\infty) \\ &= \mathbb{E}_{\lambda^*}[\mathbf{1}(\tau(x) < +\infty) \exp(-\lambda^* Z(\tau(x)))] \\ &= e^{-\lambda^* x} \mathbb{E}_{\lambda^*}[\mathbf{1}(\tau(x) < +\infty) \exp(-\lambda^* [Z(\tau(x)) - x])] \\ &\leq e^{-\lambda^* x}, \end{aligned}$$

where we used that $\mathbb{P}^{\lambda^*}(\tau(x) < +\infty) = 1$ and that the overshoot at x , $U(x) := Z(\tau(x)) - x$, is positive. In fact, if one assumes that the jump measure of Z is not supported on a lattice, then $\mathbb{P}^{\lambda^*}(\tau(x) < +\infty) \exp(\lambda^* x)$ converges to a positive constant (see Kyprianou 2014, theorem 7.6; Bertoin and Doney 1994; and Palmowski and Pistorius 2009). We will recall this fact repeatedly.

Proof of Theorems 1 and 2. We begin by noticing that for any time t (discrete or continuous),

$$\mathbb{P}\left(\tau^{(N)}(x) \leq t\right) = N\mathbb{P}\left(\tau(x) \leq t\right) + \mathbf{O}\left(\left(\frac{N}{\lceil N/2 \rceil}\right)\mathbb{P}\left(\tau(x) \leq t\right)^2\right); \quad (26)$$

so, for N fixed, $\mathbb{P}(\tau^{(N)}(x) \leq t) \sim N\mathbb{P}(\tau(x) \leq t)$. Recall that $\lceil \cdot \rceil$ denotes the ceiling function.

Case I: $N\mu(\lambda) < \mu(\lambda^*)$. In this case $B(x)$ and $B(x)/N$ are in the same regime of Proposition 1; thus, by (26),

$$\lim_{x \rightarrow +\infty} \frac{\mathbb{P}\left(\tau^{(N)}(x) \leq \frac{B(x)}{N}\right)}{\mathbb{P}\left(\tau(x) \leq B(x)\right)} = \lim_{x \rightarrow +\infty} \frac{N\mathbb{P}\left(\tau(x) \leq \frac{B(x)}{N}\right)}{\mathbb{P}\left(\tau(x) \leq B(x)\right)} = N. \quad (27)$$

For the remaining cases some considerations will be needed: for each λ such that $\mu(\lambda) > 0$, let

$$S(\lambda) := \{s \geq 1 : \exists \lambda_s > \lambda, \text{ s.t. } \mu(\lambda_s) = s\mu(\lambda)\}. \quad (28)$$

The set $S(\lambda)$ contains a (nontrivial) right neighbourhood of one as long as $\lambda < \lambda_{\max}$. We also note that for each λ as before, the map $s \mapsto \lambda_s$ is continuously differentiable and

$$\frac{d}{ds} \lambda_s = \frac{\mu(\lambda)}{\mu'(\lambda_s)} = \frac{\mu(\lambda)}{\sigma^2(\lambda_s)}.$$

Case II: $\mu(\lambda) < \mu(\lambda^*) < N\mu(\lambda)$. As a subcase, assume first that $N \in S(\lambda)$, so that there exists $\lambda_N \in (\lambda, \lambda_{\max})$ such that

$$\mu(\lambda_N) = N\mu(\lambda). \quad (29)$$

In this case, we have that $x = \mu(\lambda_N)\frac{B(x)}{N} + \mathbf{o}(\sqrt{x})$, so the exponential rate of $\mathbb{P}\left(\tau(x) \leq \frac{B(x)}{N}\right)$ is given by Proposition 1:

$$\frac{B(x)}{N} \zeta[\mu(\lambda_N)] = \frac{B(x)}{N} (\lambda_N \mu(\lambda_N) - \psi(\lambda_N)). \quad (30)$$

On the other hand, $\mathbb{P}(\tau(x) \leq B(x)) \sim e^{-\lambda^* x}$. Then, for some constant $\tilde{D}(\lambda)$,

$$\begin{aligned} \lim_{x \rightarrow +\infty} \frac{\mathbb{P}\left(\tau(x) \leq \frac{B(x)}{N}\right)}{\mathbb{P}\left(\tau(x) \leq B(x)\right)} &\leq \lim_{x \rightarrow +\infty} \tilde{D}(\lambda) \exp\left(-x(\lambda_N - \lambda^*) + \frac{B(x)}{N} \psi(\lambda_N)\right) \\ &= \lim_{x \rightarrow +\infty} \tilde{D}(\lambda) \exp\left(-x(\lambda_N - \lambda^*) \left(1 - \frac{\psi(\lambda_N) - \psi(\lambda^*)}{\mu(\lambda_N)(\lambda_N - \lambda^*)}\right) + \lambda_N \mathbf{o}(\sqrt{x})\right). \end{aligned} \quad (31)$$

Using the Mean Value Theorem and the fact that μ is strictly increasing in the interval (λ^*, λ_N) , we conclude that the limit in (31) is zero.

For the second subcase, let $N \notin S(\lambda)$, so that $N\mu(\lambda)$ is at a positive distance above the range of $\mu(\cdot)$. A bound may be obtained from any $s \in S(\lambda) \setminus \{1\}$ by monotonicity of $t \mapsto \mathbb{P}(\tau(x) \leq t)$:

$$\lim_{x \rightarrow +\infty} \frac{\mathbb{P}\left(\tau^{(N)}(x) \leq \frac{B(x)}{N}\right)}{\mathbb{P}\left(\tau(x) \leq B(x)\right)} \leq \lim_{x \rightarrow +\infty} \frac{\mathbb{P}\left(\tau^{(N)}(x) \leq \frac{x}{\mu(\lambda_s)} + \mathbf{o}(\sqrt{x})\right)}{\mathbb{P}\left(\tau(x) \leq B(x)\right)} = 0. \quad (32)$$

Case III: $\mu(\lambda^*) < \mu(\lambda)$. As in the subcases considered above, assume first that $N \in S(\lambda)$. Now both $B(x)$ and $B(x)/N$ lie in the right-hand side (RHS) of the threshold in (23). It suffices to show that the function

$$s \in S(\lambda) \mapsto \frac{1}{s} \zeta[\mu(\lambda_s)]$$

is strictly increasing. Indeed, in that case

$$\lim_{x \rightarrow +\infty} \frac{\mathbb{P}\left(\tau^{(N)}(x) \leq \frac{B(x)}{N}\right)}{\mathbb{P}\left(\tau(x) \leq B(x)\right)} = \lim_{x \rightarrow +\infty} \exp\left\{B(x) \left(\zeta[\mu(\lambda)] - \frac{1}{N} \zeta[\mu(\lambda_N)]\right)\right\} = 0, \quad (33)$$

and the result is proved.

It only remains to prove that $\frac{1}{s}\zeta[\mu(\lambda_s)]$ is increasing, and using Remark 6 this is seen to hold:

$$\begin{aligned} \frac{d}{ds} \frac{1}{s} \zeta[\mu(\lambda_s)] &= \frac{d}{ds} \lambda_s \mu(\lambda) + \frac{\psi(\lambda_s)}{s^2} - \frac{1}{s} \mu(\lambda_s) \frac{d}{ds} \lambda_s \\ &= \frac{\psi(\lambda_s)}{s^2}, \end{aligned} \tag{34}$$

which is positive because $\lambda_s > \lambda > \lambda^*$.

The proof is finished by arguing in the same way as in Case II for $N \notin S(\lambda)$. \square

5. Proof of Theorem 3

For the proof of Theorem 3 we will need the following auxiliary lemmas:

Lemma 1. *Let*

$$q(x) := \mathbb{P}_{v_x}(\tau(x) < \tau(0)). \tag{35}$$

Then under the assumptions of Theorem 3,

$$\lim_{x \rightarrow +\infty} q(x) = 0. \tag{36}$$

Moreover, there exists a constant $c > 0$ such that

$$c \int_0^x e^{\lambda^* y} v_x(dy) \leq e^{\lambda^* x} q(x) \leq \int_0^x e^{\lambda^* y} v_x(dy). \tag{37}$$

Lemma 2. *Let η and ζ be random variables defined by*

$$\mathbb{P}(\eta \leq t) = \mathbb{P}_{v_x}(\tau(0) \leq t \mid \tau(0) < \tau(x))$$

and

$$\mathbb{P}(\zeta \leq t) = \mathbb{P}_{v_x}(\tau(x) \leq t \mid \tau(x) < \tau(0)).$$

Under the assumptions of Theorem 3,

$$\limsup_{x \rightarrow +\infty} \mathbb{E}[\eta] < +\infty$$

and²

$$\mathbb{E}[\zeta] \asymp \frac{x}{\psi'(\lambda^*)}.$$

Lemma 3. *Let η, ζ be as in Lemma 2 and consider an i.i.d. sequence η_1, η_2, \dots with the same distribution as η . For a fixed $\delta > 0$, define the events*

$$E_\delta^m := \left\{ \sum_{j=1}^m \eta_j \leq m(1 + \delta) \mathbb{E}\eta_1 \right\} \text{ and } F_\delta := \{\zeta \leq (1 + \delta) \mathbb{E}\zeta\},$$

and for any $T > 0$ let

$$n_\delta(T) := \left\lfloor \frac{T - (1 + \delta) \mathbb{E}\zeta}{(1 + \delta) \mathbb{E}\eta_1} \right\rfloor.$$

In the setting of Theorem 3, the following holds:

$$\lim_{x \rightarrow +\infty} \mathbb{P}\left(E_\delta^{n_\delta(B(x))} \cap F_\delta\right) = 1. \tag{38}$$

Assuming these, we may now give the main result for restarted processes.

Proof of Theorem 3. The first time the restarted process goes above x has a regenerative structure with cycles determined by the restart times. Formally,

$$\tau^{v_x}(x) = \sum_{j=1}^{N_x} \eta_j + \zeta, \tag{39}$$

where $N_x \in \{0, 1, \dots\}$ is defined as the number of “failed” cycles (i.e., those ending at time $\tau(0)$) before a successful one (i.e., one ending at time $\tau(x)$) and has a geometric distribution with success parameter $q(x) := \mathbb{P}_{v_x}(\tau(x) < \tau(0))$. For each $x > 0$ the laws of the η_j 's and of ζ are those of $\tau(0)$ conditioned on the event $\{\tau(0) < \tau(x)\}$ and of $\tau(x)$ conditioned on $\{\tau(x) < \tau(0)\}$, respectively. We remark that although it is not explicit in our notation, the laws of η_1 and ζ both depend on x .

Lemma 1 studies the behavior of $q(x)$ for large x ; in particular (36) shows that it becomes increasingly harder to observe the events of interest as x grows.

A lower bound for $\mathbb{P}(\tau^{v_x}(x) \leq B(x))$ is given by the probability that N_x in (39) takes an unusually small value. To this end, consider for a fixed small $\delta > 0$ the events

$$E_\delta^m := \left\{ \sum_{j=1}^m \eta_j \leq m(1 + \delta) \mathbb{E}\eta_1 \right\} \quad \text{and} \quad F_\delta := \{\zeta \leq (1 + \delta) \mathbb{E}\zeta\},$$

and for each $T > 0$,

$$n_\delta(T) := \lfloor \frac{T - (1 + \delta) \mathbb{E}\zeta}{(1 + \delta) \mathbb{E}\eta_1} \rfloor.$$

Then by Lemma 3,

$$\begin{aligned} \liminf_{x \rightarrow +\infty} \mathbb{P}(\tau^{v_x}(x) \leq B(x)) &\geq \liminf_{x \rightarrow +\infty} \mathbb{P}(\{N_x \leq n_\delta(B(x))\} \cap E_\delta^{n_\delta(B(x))} \cap F_\delta) \\ &= \liminf_{x \rightarrow +\infty} \mathbb{P}(N_x \leq n_\delta(B(x))). \end{aligned}$$

Finally, we can conclude the desired uniform lower bound for (15) because

$$\mathbb{P}(N_x \leq n_\delta(B(x))) = 1 - (1 - q(x))^{n_\delta(B(x))} = n_\delta(B(x)) q(x) (1 + o(1))$$

and $n_\delta(B(x))/B(x)$ is bounded away from zero and infinity by Lemma 2. It follows that for $B(x)$ in the regime considered

$$\mathbb{P}(\tau(x) \leq B(x)) B(x) q(x) \asymp \mathbb{P}(\tau(x) \leq B(x)) B(x) \int_0^x \exp(\lambda^* y) v_x(dy).$$

The proof of Theorem 3 is now concluded. \square

5.1. Proofs of Auxiliary Lemmas

Proof of Lemma 1. We begin with the proof of the upper bound of (37) by a martingale argument. We have

$$\begin{aligned} \int_0^x e^{\lambda^* y} v_x(dy) &= \mathbb{E}_{v_x} M_0 = \mathbb{E}_{v_x} M_{\tau(x) \wedge \tau(0)} \\ &= e^{\lambda^* x} \mathbb{E}_{v_x} [e^{\lambda^*(Z(\tau(x)) - x)} \mathbf{1}(\tau(x) < \tau(0))] + \mathbb{E}_{v_x} [e^{\lambda^* Z(\tau(0))} \mathbf{1}(\tau(x) > \tau(0))] \\ &\geq e^{\lambda^* x} \mathbb{E}_{v_x} [\mathbf{1}(\tau(x) < \tau(0))] = e^{\lambda^* x} q(x). \end{aligned}$$

We now turn to the lower bound. Our main ingredient is Bertoin and Doney (1994),³ who state that there exists a constant $C \in (0, 1)$ such that

$$\lim_{x \rightarrow +\infty} e^{\lambda^* x} \mathbb{P}(\tau(x) < +\infty) = C. \tag{40}$$

Using that $\tau(0)$ is almost surely finite, for every x and $y \in (0, x)$

$$\mathbb{P}_y(\tau(x) < +\infty) = \mathbb{P}_y(\tau(0) < \tau(x) < +\infty) + \mathbb{P}_y(\tau(x) < \tau(0)). \tag{41}$$

We first claim that

$$l = \liminf_{x \rightarrow +\infty} \int_0^x e^{\lambda^* y} v_x(dy) > 1. \tag{42}$$

Indeed, $l \geq 1$ and equality cannot hold because otherwise the measures v_x would converge (up to taking a subsequence) to the delta measure on zero, thus contradicting the positive limit of the absorption rates in Lemma 6, and the claim follows.

In view of the claim above, let $\delta > 0$ be small enough such that $C(1 - \delta) > C \frac{(1+\delta)}{l}$, and let also x_δ be given by (40) such that for $x > x_\delta$

$$C(1 - \delta) \leq e^{\lambda^* x} \mathbb{P}(\tau(x) < +\infty) \leq C(1 + \delta).$$

Noticing that by the Markov property $\mathbb{P}_y(\tau(0) < \tau(x) < +\infty) \leq \mathbb{P}_0(\tau(x) < +\infty) = \mathbb{P}(\tau(x) < +\infty)$, we then have

$$\begin{aligned} q(x) &= \int_0^x \mathbb{P}_y(\tau(x) < \tau(0)) \nu_x(dy) \\ &= \int_0^x \mathbb{P}(\tau(x-y) < +\infty) - \mathbb{P}_y(\tau(0) < \tau(x) < +\infty) \nu_x(dy) \\ &\geq \int_0^x \mathbb{P}(\tau(x-y) < +\infty) - \mathbb{P}(\tau(x) < +\infty) \nu_x(dy) \\ &\geq C(1 - \delta) e^{-\lambda^* x} \int_0^{x-x_\delta} e^{\lambda^* y} \nu_x(dy) + \int_{x-x_\delta}^x \mathbb{P}(\tau(x-y) < +\infty) \nu_x(dy) - C(1 + \delta) e^{-\lambda^* x} \\ &\geq C(1 - \delta) e^{-\lambda^* x} \int_0^{x-x_\delta} e^{\lambda^* y} \nu_x(dy) + \mathbb{P}(\tau(x_\delta) < +\infty) \nu_x([x - x_\delta, x]) - C(1 + \delta) e^{-\lambda^* x}. \end{aligned} \tag{43}$$

It then suffices to find a strictly positive lower bound, independent of x for the following:

$$\begin{aligned} \frac{q(x) e^{\lambda^* x}}{\int_0^x e^{\lambda^* y} \nu_x(dy)} &\geq C(1 - \delta) \frac{\int_0^{x-x_\delta} e^{\lambda^* y} \nu_x(dy)}{\int_0^x e^{\lambda^* y} \nu_x(dy)} \\ &\quad + \mathbb{P}(\tau(x_\delta) < +\infty) \nu_x([x - x_\delta, x]) \frac{e^{\lambda^* x}}{\int_0^x e^{\lambda^* y} \nu_x(dy)} \\ &\quad - C(1 + \delta) \frac{1}{\int_0^x e^{\lambda^* y} \nu_x(dy)}. \end{aligned} \tag{44}$$

There are two cases to distinguish: whether $l(\delta) = \liminf_{x \rightarrow +\infty} \frac{\int_0^{x-x_\delta} e^{\lambda^* y} \nu_x(dy)}{\int_0^x e^{\lambda^* y} \nu_x(dy)}$ is equal to one or not.

In the case $l(\delta) = 1$, the first term in the RHS of (44) dominates the third and the desired lower bound is $C(1 - \delta) - C(1 + \delta)/l > 0$.

If we now assume that $l(\delta) < 1 - \varepsilon$ for some positive ε , then the middle term in (44) is at least $\varepsilon \mathbb{P}(\tau(x_\delta) < +\infty)$ for large x . The proof concludes by showing that either the first and third terms vanish or the first dominates the third: assume first that there exists a sequence $x_n \rightarrow +\infty$ such that

$$\int_0^{x_n - x_\delta} e^{\lambda^* y} \nu_{x_n}(dy) < 1 - \varepsilon'$$

for some $\varepsilon' > 0$ and all n ; then $\nu_{x_n}([x_n - x_\delta, x_n]) \geq \varepsilon'$ and consequently

$$\int_0^{x_n} e^{\lambda^* y} \nu_{x_n}(dy) \rightarrow +\infty.$$

For such a sequence, the first and third terms on the RHS of (44) vanish as $x \rightarrow +\infty$. On the other hand, if there is no such sequence (and hence $\liminf_{x \rightarrow +\infty} \int_0^{x-x_\delta} e^{\lambda^* y} \nu_x(dy) \geq 1$), the first term dominates the third one and $\varepsilon \mathbb{P}(\tau(x_\delta) < +\infty)$ is still a lower bound for (44).

Let us now observe that $q(x)$ goes to zero as $x \rightarrow +\infty$, so the events of interest are indeed rare events: for any $\varepsilon > 0$, there is k_ε such that $\nu[k_\varepsilon, +\infty) < \varepsilon$; then for x big enough, we have that

$$\begin{aligned} q(x) &\asymp \exp(-\lambda^* x) \int_0^x \exp(\lambda^* y) \nu_x(dy) \\ &\leq \exp(-\lambda^* x) \int_0^{k_\varepsilon} \exp(\lambda^* y) \nu_x(dy) + \nu_x[k_\varepsilon, x]. \end{aligned}$$

Because k_ε is fixed, the above limit is upper bounded by ε , and ε is arbitrary. Hence, $q(x) \rightarrow 0$ as $x \rightarrow +\infty$.

The proof of Lemma 1 is then complete. \square

Proof of Lemma 2. The statement on η follows easily from stochastic domination by noticing that the function $y \mapsto \mathbb{E}_y \tau(0)$ is nondecreasing. Indeed,

$$\limsup_{x \rightarrow +\infty} \mathbb{E}[\eta] = \limsup_{x \rightarrow +\infty} \frac{\mathbb{E}_{\nu_x}[\tau(0) \mathbf{1}(\tau(0) < \tau(x))]}{1 - q(x)} \leq \frac{\mathbb{E}_\nu[\tau(0)]}{1 + \mathbf{o}(1)} < +\infty. \tag{45}$$

We now turn to ζ . For each $x > 0$ denote by $U(x)$ the overshoot at x : $U(x) = Z(\tau(x)) - x$.

by means of the exponential change of measure and the asymptotics of

$$\begin{aligned} \lim_x \mathbb{E}[\zeta] &= \lim_x \mathbb{E}_{v_x}[\tau(x) \mid \tau(x) < \tau(0)] \\ &= \lim_x \frac{\mathbb{E}_{v_x}^{\lambda^*}[e^{-\lambda^* U(x) + \lambda^* y_0} ; \tau(x) < \tau(0)]}{\mathbb{E}_{v_x}^{\lambda^*}[e^{-\lambda^* U(x) + \lambda^* y_0} ; \tau(x) < \tau(0)]} \\ &\asymp \mathbb{E}_{v_x}^{\lambda^*}[\tau(x)] \\ &\asymp \frac{1}{\psi'(\lambda^*)} \int_0^x (x - y_0) v_x(dy_0). \end{aligned}$$

Using that v has finite first moment, the last term above is seen to scale as $x/\psi'(\lambda^*)$ up to multiplicative constants. \square

Proof of Lemma 3. To lighten the notation, let us denote $n_\delta(B(x))$ by n_δ . We will show separately that

$$\lim_x \mathbb{P} \left(\frac{1}{n_\delta} \sum_{j=1}^{n_\delta} \eta_j \geq (1 + \delta) \mathbb{E}\eta_1 \right) = 0 \quad (46)$$

and

$$\lim_x \mathbb{P}(\zeta > (1 + \delta) \mathbb{E}\zeta) = 0. \quad (47)$$

Together they imply (38).

Let us start with (46). For this part of the proof, we apply Chebyshev's inequality:

$$\begin{aligned} \lim_x \mathbb{P} \left(\frac{1}{n_\delta} \sum_{j=1}^{n_\delta} \eta_j \geq (1 + \delta) \mathbb{E}\eta_1 \right) &\leq \lim_x \mathbb{P} \left(\frac{1}{n_\delta} \sum_{j=1}^{n_\delta} |\eta_j - \mathbb{E}\eta_1|^2 \geq (\delta \mathbb{E}\eta_1)^2 \right) \\ &\leq \frac{\text{Var}(\eta_1)}{n_\delta (\delta \mathbb{E}\eta_1)^2} \\ &\leq \frac{1}{n_\delta (\delta \mathbb{E}\eta_1)^2} \int_0^{+\infty} \mathbb{E}_y[\tau(0)^2] v(dy) \\ &\asymp \frac{1}{n_\delta (\delta \mathbb{E}\eta_1)^2} \int_0^{+\infty} (x - y)^2 v(dy). \end{aligned}$$

The asymptotic equivalence of the last line above is justified by Doney and Maller (2004, theorem 3) once we note that $\mathbb{P}(Z(t) \geq -x_0)$ decays exponentially to zero with t for every fixed $x_0 > 0$ (this is easily checked with an exponential martingale argument using any parameter $\theta > 0$ such that $\psi(\theta) < 0$).

Because v has finite second moment and does not equal δ_0 , $\mathbb{E}[\eta_1]$ does not go to zero with $x \rightarrow +\infty$ and we conclude (46).

Let us now prove (47). On the one hand, the Strong Law of Large Numbers (see Kyprianou 2014, exercise 7.2, p. 224) implies that

$$\lim_{x \rightarrow +\infty} \frac{Z(\tau(x))}{\tau(x) \mathbb{E}^{\lambda^*} Z(1)} = \lim_{x \rightarrow +\infty} \frac{x + U(x)}{\tau(x) \psi'(\lambda^*)} = 1, \mathbb{P}^{\lambda^*} - a.s.$$

Here as before $U(x)$ denotes the overshoot at x . On the other hand, as seen in Lemma 2, for large x , $\mathbb{E}\zeta$ is asymptotically equivalent to $x/\psi'(\lambda^*)$. Under Cramér's condition $U(x)/x$ converges to zero almost surely, so we conclude that

$$\lim_{x \rightarrow +\infty} \frac{\tau(x)}{\mathbb{E}^{\lambda^*}[\tau(x)]} = 1, \mathbb{P}^{\lambda^*} - a.s.$$

and (47) follows. \square

6. Proof of Theorem 4

In this section Z denotes either a random walk on discrete time or a Lévy process, satisfying Cramér's condition, and we will consider as restarting measures the family of quasi-stationary distributions $(v_x)_{x \geq 1}$. The proof of Theorem 4 relies on some auxiliary results.

Lemma 4. The sequence of QSD measures $(\nu_x)_{x \geq 1}$ converges in distribution to the Yaglom limit on $(0, +\infty)$, given by

$$\nu(B) := \lim_{t \rightarrow +\infty} \mathbb{P}_{x_0}(Z(t) \in B \mid t < \tau(0)),$$

where $\tau(0) = \inf\{t > 0 : Z(t) \leq 0\}$ and $x_0 \in (0, +\infty)$ is an arbitrary point.

Once the sequence of measures admits a weak limit, Remark 2 provides a route to apply Theorem 3, thereby yielding the lower bound in Theorem 4. We adopt, however, a different approach that relies more heavily on properties associated with quasi-stationarity.

Lemma 5. For the restarted process Z^{ν_x} , $\tau^{\nu_x}(x)$ has an exponential distribution with parameter $\beta(x)q(x)$, where $\beta(x)$ is the exponential rate of absorption associated with ν_x and $q(x)$ is defined in (35).

Lemma 6. The function $\beta(x)$, as defined in Lemma 5, is decreasing in $x > 0$ and has a positive limit $\beta > 0$ as $x \rightarrow +\infty$.

If we assume the lemmas above, we may prove the main result as follows:

Proof of Theorem 4. By Lemma 5 we have that $\mathbb{P}(\tau^{\nu_x}(x) \leq B(x)) = 1 - \exp[-\beta(x)q(x)B(x)]$. Because $B(x)$ grows slower than e^{λ^*x} with x , $\beta(x)q(x)B(x) \rightarrow 0$ and $\mathbb{P}(\tau^{\nu_x}(x) \leq B(x)) \sim \beta(x)q(x)B(x)$ as $x \rightarrow +\infty$, with the notation \sim defined at the beginning of Subsection 4.2.

Also, by Lemma 6, $\beta(x)$ has a positive limit, and because $\mathbb{P}(\tau(x) \leq B(x))$ has order $\exp(-\lambda^*x)$, by Proposition 1 we conclude that the ratio

$$\frac{\mathbb{P}(\tau^{\nu_x}(x) < B(x))}{\mathbb{P}(\tau(x) < B(x))} \frac{1}{B(x) \int_0^x e^{\lambda^*y} \nu_x(dy)} \quad (48)$$

is bounded away from zero and infinity uniformly in x . \square

Before proving the lemmas, we compute the quasi-stationary distribution, its exponential moment of order λ^* , and the absorption rate for a linear Brownian motion explicitly.

Proof of Corollary 2. In the case in which Z is distributed as a Brownian motion, we can compute all ingredients explicitly. Brownian motion with constant drift $-\mu < 0$ has infinitesimal generator

$$\mathcal{L} = \frac{1}{2} \frac{d^2}{dx^2} - \mu \frac{d}{dx},$$

whose domain contains the set of twice continuously differentiable functions with right limit (resp. left limit) equal to zero at zero (resp. x), denoted by $C_0^2(0, x)$. The adjoint operator is

$$\mathcal{L}^* = \frac{1}{2} \frac{d^2}{dx^2} + \mu \frac{d}{dx}$$

and has $u(y) = \sin(\pi y/x)e^{-\mu y}$ as a positive eigenfunction. Indeed,

$$\mathcal{L}^* u = -\frac{1}{2} \left(\mu^2 + \frac{\pi^2}{x^2} \right) u.$$

By the spectral characterization of quasi-stationary measures (see, e.g., Méléard and Villemonais 2012, proposition 4), the measure

$$\nu_x(dy) = D \sin(\pi y/x) e^{-\mu y} dy \quad (49)$$

is then quasi-stationary for Z , where $D = \frac{\mu^2 x^2 + \pi^2}{\pi x(e^{-\mu x} + 1)}$ is the normalization constant. The absorption time $T_x = \tau(x) \wedge \tau(0)$ is exponentially distributed with rate $\beta(x) = \frac{1}{2} \left(\mu^2 + \frac{\pi^2}{x^2} \right)$, that is,

$$\mathbb{P}_{\nu_x}(T_x > t) = e^{-\frac{1}{2} \left(\mu^2 + \frac{\pi^2}{x^2} \right) t}. \quad (50)$$

With the explicit form for the quasi-stationary distribution, and recalling that $\lambda^* = 2\mu$, the exponential moment follows:

$$\int_0^x e^{\lambda^* y} \nu_x(dy) = D \int_0^x e^{\mu y} \sin\left(\frac{\pi}{x} y\right) dy = D \pi x \frac{e^{\mu x} + 1}{\mu^2 x^2 + \pi^2} = e^{\mu x}. \quad (51)$$

The result now follows from Theorem 4. \square

6.1. Proofs of Auxiliary Lemmas

We now turn to the proof of the auxiliary lemmas.

Proof of Lemma 4. Let us recall the piece of notation $T_x = \inf\{t > 0 : Z(t) \notin (0, x)\} = \tau(0) \wedge \tau(x)$ and take x_0 any number in $(0, +\infty)$. We shall prove that the following limit holds for any bounded measurable function f :

$$\lim_{x \rightarrow +\infty} \lim_{t \rightarrow +\infty} |\mathbb{E}_{x_0}[f(Z(t)) | t < \tau(0)] - \mathbb{E}_{x_0}[f(Z(t)) | t < T_x]| = 0. \quad (52)$$

The result relies on proving that $\mathbb{P}_{x_0}(t < \tau(x) \wedge \tau(0)) \sim \mathbb{P}_{x_0}(t < \tau(0))$ when taking t and x to $+\infty$ in the same order as in (52). To check this fact, we start by using the strong Markov property

$$\begin{aligned} 1 - \frac{\mathbb{P}_{x_0}(t < T_x)}{\mathbb{P}_{x_0}(t < \tau(0))} &= \frac{\mathbb{P}_{x_0}(\tau(x) \leq t < \tau(0))}{\mathbb{P}_{x_0}(t < \tau(0))} \\ &= \frac{\mathbb{E}_{x_0}\{\mathbf{1}(\tau(x) \leq t) \mathbb{P}_{Z(\tau(x))}(t - \tau(x) < \tau(0))\}}{\mathbb{P}_{x_0}(t < \tau(0))}. \end{aligned} \quad (53)$$

In Denisov and Shneer (2013, theorem 3.5), it is proved that for Lévy processes, for large t ,

$$\mathbb{P}_{x_0}(t < \tau(0)) \sim \text{const } V(x_0) \exp(t \psi(\lambda_0)) t^{-3/2},$$

where $\lambda_0 \in (0, \lambda^*)$ is the unique positive solution to $\psi'(\lambda) = 0$, and the function V as defined in Denisov and Shneer (2013, theorem 2.1) satisfies

$$g(x) e^{\lambda_0 x} \leq V(x) \leq e^{\lambda_0 x} C \quad (54)$$

for some constant $C > 0$ and some bounded increasing function g . It now follows that the right-hand side of (53) converges to zero. Indeed,

$$\begin{aligned} 1 - \frac{\mathbb{P}_{x_0}(t < T_x)}{\mathbb{P}_{x_0}(t < \tau(0))} &\leq \text{const} \frac{\mathbb{E}_{x_0}\{\mathbf{1}(\tau(x) \leq t) V(Z(\tau(x)))\}}{V(x_0)} \\ &\stackrel{t \rightarrow +\infty}{\leq} \text{const} \frac{1}{\exp(\lambda_0 x_0)} \mathbb{E}_{x_0}\{\mathbf{1}(\tau(x) \leq +\infty) \exp\{\lambda_0(x - x_0 + U(x))\}\} \\ &\leq \text{const} \mathbb{E}^{\lambda^*}\{\exp(-(\lambda^* - \lambda_0)(x - x_0 + U(x)))\} \\ &= \mathbf{o}(1), x \rightarrow +\infty. \end{aligned}$$

Here we used the exponential martingale change of measure and the overshoot $U(x)$ once more (recall Remark 7). In the case of a random walk under Cramér's condition, Denisov and Shneer (2013, theorem 3.5) provide an analogous estimate but with a spatial factor $V = 1$, so all the above considerations still hold. We may now conclude (52):

$$\begin{aligned} &|\mathbb{E}_{x_0}[f(Z(t)) | t < \tau(0)] - \mathbb{E}_{x_0}[f(Z(t)) | t < T_x]| \\ &= \frac{|\mathbb{E}_{x_0}[f(Z(t)); t < \tau(0)] - (1 + \mathbf{o}(1)) \mathbb{E}_{x_0}[f(Z(t)); t < T_x]|}{\mathbb{P}_{x_0}(t < \tau(0))} \\ &\leq \frac{|\mathbb{E}_{x_0}[f(Z(t)); \tau(x) \leq t < \tau(0)]|}{\mathbb{P}_{x_0}(t < \tau(0))} + \mathbf{o}(1) \frac{|\mathbb{E}_{x_0}[f(Z(t)); t < T_x]|}{\mathbb{P}_{x_0}(t < \tau(0))} \\ &\leq \|f\|_\infty \mathbf{o}(1), \end{aligned} \quad (55)$$

and the proof is complete. \square

Proof of Lemma 5. We recall that for the process Z absorbed upon exiting $(0, x)$, the absorption time $\tau(x) \wedge \tau(0)$ is exponentially distributed with parameter $\beta(x)$ if started with the quasi-stationary distribution ν_x (see, e.g., Collet et al. 2013, theorem 2.2). Observe that the $\tau^{\nu_x}(x)$ can be decomposed into a sum of a geometric number of duration of cycles, namely,

$$\tau^{\nu_x}(x) = \sum_{j=1}^{N_x} \tau_j, \quad (56)$$

where N_x is the index of the first cycle at which Z exits through $[x, +\infty)$, and τ_j is the duration of the j -th cycle. Thus, N_x is geometric with parameter $q(x) = \mathbb{P}_{\nu_x}(\tau(x) < \tau(0))$ and the times τ_j are independent and distributed as $\tau(x) \wedge \tau(0)$. Then $\tau^{\nu_x}(x)$ is a geometric sum of i.i.d. exponential random variables. Moreover, because ν_x is a quasi-

stationary distribution, the time and position at which Z exits $(0, x)$ are independent (Collet et al. 2013, theorem 2.6), and hence so is N_x of $\tau_1, \dots, \tau_{N_x}$. Being a geometric sum of exponential random variables which are mutually independent and independent of the number of terms, we conclude that $\tau^{v_x}(x)$ has an exponential distribution of parameter $\beta(x)q(x)$. \square

Proof of Lemma 6. Let $T_x = \tau(x) \wedge \tau(0) = \inf\{t > 0 : Z(t) \notin (0, x)\}$ be the first exit time from $(0, x)$; because the quasi-stationary measure is the unique Yaglom limit for the system under consideration, the rate of absorption $\beta(x)$ is obtained as

$$\beta(x) = - \lim_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{P}_y(T_x > t) \quad (57)$$

for any $y \in (0, x)$. Because $x' > x > y > 0$ implies $T_{x'} \geq T_x$ \mathbb{P}_y -almost surely, it follows that $\beta(x') \geq \beta(x)$. The sequence of absorption rates is nonincreasing with x , and hence has a limit β . The same reasoning shows that if β_∞ denotes the absorption rate of the QSD on $(0, +\infty)$, then $\beta(x) \geq \beta_\infty$ for every $x \geq 1$, which implies that $\beta \geq \beta_\infty > 0$. \square

7. Conclusion

This paper investigates various exploration strategies under time constraints in environments with unknown stochastic dynamics focusing on their impact on performance as measured by the time required to reach a set of rare states.

We aim for this work to be a foundational step toward developing a more qualitative theory of exploration, specifically by incorporating the time needed to observe meaningful signals for the first time. For example, in Reinforcement Learning environments with sparse rewards, exploration of the state space or action-state space is widely recognized as a critical bottleneck to efficiency. By addressing this crucial aspect, we seek to contribute to more efficient and effective exploration strategies in such contexts and beyond.

In this work, we focused on space-invariant one-dimensional dynamics, which provide highly interpretable results and for which we can provide explicit performance guarantees. Building on this understanding, future work will consider more general Markovian dynamics, providing a more comprehensive and realistic framework for exploring challenging environments. A natural extension of this work is to higher-dimensional settings. Although explicit hitting-time estimates of the kind used here are generally unavailable in that regime, the “too-many-particles” phenomenon and the core conclusions of our main results are expected to carry over to considerably more general settings.

Acknowledgments

The authors are grateful to the anonymous referees for their thorough and insightful reviews; their suggestions and observations led to substantial improvements in this work. P.B. and E.G. also express their deep gratitude to professors E. Mordecki and J. R. León for very valuable discussions.

Appendix A. M/M/1

In this section we investigate the interplay between the probability of exploration, under the time regime of interest, and the efficiency of the associated estimation procedure. In particular, we show that the variance of the natural Monte Carlo estimator mirrors the typical exploration time.

Consider an $M/M/1$ queue on the truncated state space $\{0, 1, \dots, n\}$ with load parameter $\rho < 1$, and assume that $X(0) = 0$ almost surely. The chain is ergodic, and its invariant measure is geometrically decaying: for any $k \in \mathbb{N}$ with $k \leq n$, the stationary weight is proportional to ρ^k . By the Ergodic Theorem, the stationary measure π satisfies

$$\pi(k) = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \mathbf{1}(X(s) = k) ds.$$

A Monte Carlo estimator for $\pi(k)$, based on a time horizon $B(k)$, is therefore

$$\hat{\pi}(k) = \frac{1}{B(k)} \int_0^{B(k)} \mathbf{1}(X(s) = k) ds.$$

We focus on restrictive time budgets and assume that the estimation of $\pi(k)$ for large k is performed with a linear-in- k horizon,

$$B(k) \asymp k, \quad k \rightarrow \infty.$$

In such a regime, the process is expected to spend exceedingly little time in state k : indeed, the expected hitting time of k grows exponentially fast in k . As we show next, this scarcity of observations is the principal obstruction to reliable estimation, and it directly governs the variance of the Monte Carlo estimator.

Fix $c > 0$ and introduce the surrogate estimator

$$\xi(k) := c \mathbf{1}(\tau(k) < B(k)),$$

where $\tau(k)$ denotes the hitting time of state k . Then, by noting that $\hat{\pi}(k) + \xi(k)$ is almost surely bounded by $1 + c$, and that $|\hat{\pi}(k) - \xi(k)| \mathbf{1}(\tau(k) > B(k)) = 0$, one obtains

$$\begin{aligned} |\text{Var } \hat{\pi}(k) - \text{Var } \xi(k)| &= |\mathbb{E}_0[(\hat{\pi}(k) - \xi(k))(\hat{\pi}(k) + \xi(k))] \\ &\quad - \mathbb{E}_0(\hat{\pi}(k) - \xi(k)) \mathbb{E}_0(\hat{\pi}(k) + \xi(k))| \\ &\leq 2(1 + c) \cdot \mathbb{E}_0 |\hat{\pi}(k) - \xi(k)| \\ &\leq \text{const} \cdot \mathbb{P}_0(\tau(k) \leq B(k)), \end{aligned}$$

for some constant independent of k , thus showing that the variance is essentially governed by the rare-event probability of reaching k within the available time budget.

Appendix B. The Fleming-Viot (FV) Particle System and FVRL

The approach found in Mastropietro et al. (2025) for defining a restarted mechanism is based on the paradigm of the FV particle system, introduced by Burdzy et al. (1996). In this framework, a population of particles (which can represent different simulations in parallel) evolves over time, exploring the state space in parallel. When a particle falls into a less promising region A , it is “restarted” by being replaced with a copy of another, more promising particle. This ensures that resources are concentrated on exploring the most fruitful areas of the state space. The FV strategy is particularly effective in scenarios where certain regions of the state space are more likely to yield valuable rewards, as it dynamically reallocates exploration effort toward these regions, thereby increasing the overall efficiency of the exploration process.

Note that various papers (Ferrari and Marić 2007, Asselah et al. 2011, Villemonais 2011) have shown that FV empirical measures converge as the number of particles tends to infinity to the conditioned evolution of the process, that is,

$$m_t^N(A) \rightarrow P(X_t \in A | T \geq t), N \rightarrow \infty,$$

where T is the hitting time of A . On the other hand, $P(X_t \in A | T \geq t) \rightarrow \nu(A)$, as $t \rightarrow \infty$ where ν is the QSD (a QSD, for countable state space) associated with the process absorbed in A . This is the motivation for restarting according to the QSD in our work.

Appendix C. An Example in RL: The Control of Blocking in an $M/M/1/K$ Queue

This example has been considered in Mastropietro et al. (2025) and generalized to multidimensional settings (several coupled queues). It serves here as a canonical illustration of Reinforcement Learning in environments with *sparse and rare rewards* where the exploration is the main bottleneck. We summarize the model and the FVRL strategy introduced in Mastropietro et al. (2025) as a motivation for our theoretical results. Consider an $M/M/1/K$ queue with fixed buffer capacity K , arrival rate λ , service rate μ , and load $\rho = \lambda/\mu < 1$. The state space is $\mathcal{S} = \{0, 1, \dots, K\}$, representing queue occupancy. The control action $a \in \{0, 1\}$ determines whether to accept ($a = 1$) or block ($a = 0$) arriving jobs. The reward function is structured to penalize both the underutilisation of resource and large blocking probabilities

$$r(x, a) = \begin{cases} 0 & \text{if } a = 1 \\ B(1 + b^{x-x_{\text{ref}}}) & \text{if } a = 0 \end{cases}$$

with $B > 0$, $b > 1$, and reference state x_{ref} . This structure induces threshold-optimal policies (blocking all the incoming traffic after a given threshold). For threshold policies parameterized by θ , we define the policy-dependent blocking threshold $K_\theta = \lfloor \theta \rfloor + 1$, representing the state at which the policy π_θ begins rejecting incoming jobs. To ensure that the blocking threshold never exceeds the physical buffer, the policy parameter is restricted to the compact domain $\Theta = [0, K - 1]$, so that $K_\theta \leq K$ is guaranteed at all times. Under the average reward criterion and using a policy gradient strategy, one aims at minimizing

$$J^{\pi_\theta} = \sum_{x \in \mathcal{S}} p^{\pi_\theta}(x) \sum_a \pi_\theta(a|x) Q_\theta(x, a),$$

where $\theta \in \Theta$ is the parameterization of the policy. The policy gradient theorem yields

$$\nabla_\theta J^{\pi_\theta} = \sum_{x \in \mathcal{S}} p^{\pi_\theta}(x) \sum_a Q_\theta(x, a) \nabla_\theta \pi_\theta(a|x). \quad (\text{C.1})$$

For threshold policies parameterized by $\theta \in \Theta$, this simplifies to

$$\nabla_\theta J^{\pi_\theta} \propto p^{\pi_\theta}(\mathbf{K}_\theta - 1) [Q_\theta(\mathbf{K}_\theta - 1, 1) - Q_\theta(\mathbf{K}_\theta - 1, 0)], \quad (\text{C.2})$$

where $K_\theta = \lfloor \theta \rfloor + 1$ is the learned blocking threshold, and the gradient is concentrated at state $K_\theta - 1 = \lfloor \theta \rfloor$, the last state before the policy blocks under the current parameterization $\theta \in \Theta$.

C.1. The Gradient Estimation Problem

Given this framework, the RL strategy will be efficient if the gradient estimates can be informative. We discuss here only the estimate of the stationary probability $p^{\pi_\theta}(\mathbf{K}_\theta - 1)$. The stationary distribution for state \mathbf{K}_θ , under the threshold policy π_θ , is given by the $M/M/1/K_\theta$ formula:

$$p^{\pi_\theta}(\mathbf{K}_\theta) = \frac{(1 - \rho)\rho^{\mathbf{K}_\theta}}{1 - \rho^{\mathbf{K}_\theta+1}}.$$

For typical parameters ($\rho = 0.7$, $K_\theta = 40$), $p^{\pi_\theta}(\mathbf{K}_\theta) \sim \mathcal{O}(10^{-7})$. Because $p^{\pi_\theta}(\mathbf{K}_\theta - 1) \sim p^{\pi_\theta}(\mathbf{K}_\theta)/\rho$, both are exponentially small.

Consequently,

- Vanilla Monte Carlo (MC) estimators of $p^{\pi_\theta}(\mathbf{K}_\theta - 1)$ exhibit prohibitive variance (in terms of multiplicative errors);
- Much more importantly, the policy gradient in (C.2) becomes numerically zero under finite sampling. Hence, the learning signal vanishes, preventing convergence to an optimal θ .

C.2. Parallel Sampling

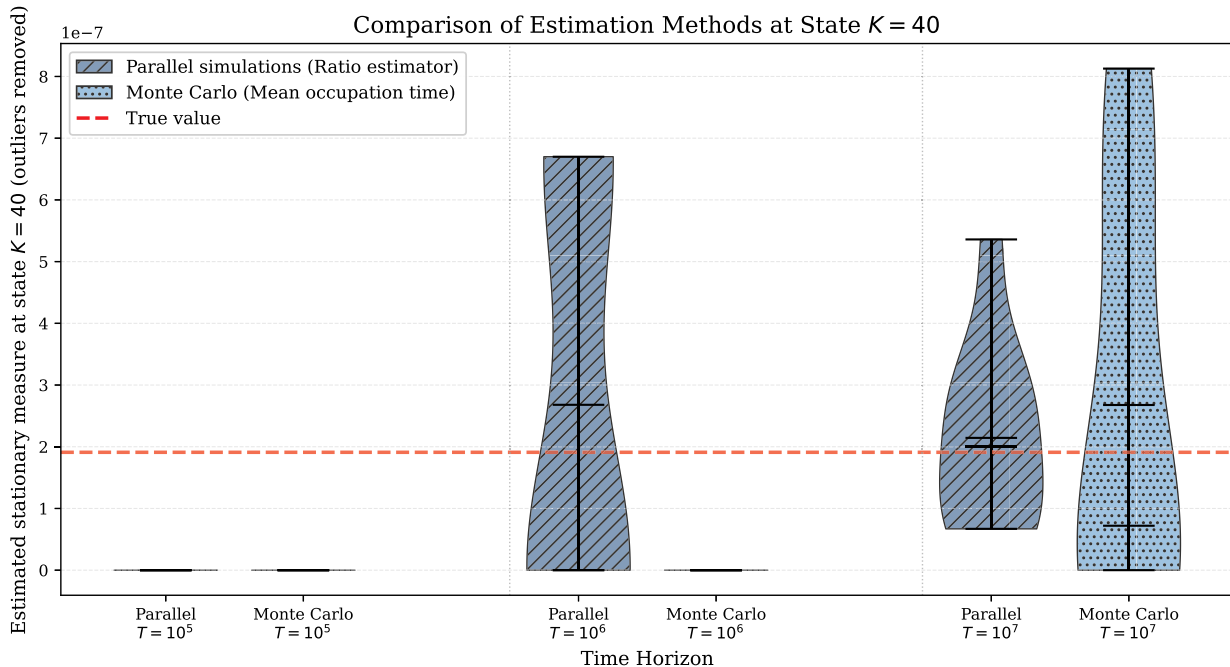
In order to compare our results with those of Mastropietro et al. (2025), we present a set of simulations for parallel $M/M/1$ queues under different time regimes. The hyperparameters used in Figure C.1 are total queue capacity $K = 40$, initial condition $J = 12$, arrival rate $\lambda = 0.7$, and service rate $\mu = 1$. In this numerical illustration, $K_\theta = K = 40$, so that $p^{\pi_\theta}(K)$ is the exponentially rare quantity identified in Subsection C.1.

Following Mastropietro et al. (2025), we estimate the stationary probability via a renewal-type decomposition based on successive returns to J . More precisely, we estimate (i) the expected return time to state J , which in our setting is of order 10^2 , and independently (ii) the probability of hitting K starting from J . For the latter we report estimators under time horizons $B(K)$ of magnitudes 10^5 , 10^6 , and 10^7 , and we compare them with the true stationary value $\mathbf{p}^{\pi_\theta}(\mathbf{K})$ shown in Figure C.1.

To ensure a fair comparison with the experiment of Mastropietro et al. (2025, section 4.1), we remark that the total number of events observed in our simulations—counting both arrivals and departures—is one to two orders of magnitude larger than theirs. Indeed, across our time regimes we record approximately 1.7×10^7 events for $r = 5, 6, 7$, which exceeds the event counts reported in Mastropietro et al. (2025, figure 2). This discrepancy is natural: correlations in the Fleming-Viot particle system substantially facilitate the exploration of the rare state K .

In summary, and as expected, the estimation accuracy with independent particles lies between that of naive Monte Carlo (which rarely reaches K even for time of order 10^6) and that of the Fleming-Viot particle system.

Figure C.1. Simulation of Parallel $M/M/1$ Queues



Notes. Simulation of parallel $M/M/1$ queues with $K = 40$, $J = 12$, $\lambda = 0.7$, and $\mu = 1$, comparing renewal-based stationary probability estimates across time horizons 10^5 , 10^6 , and 10^7 . The number of parallel copies to deploy in each time regime is computed using the results of our main theorems.

C.3. Fleming-Viot Solution

The FVRL method addresses this by introducing an *absorption set*, where prior knowledge is used on the fact that no rewards are granted in $\mathcal{A} = \{0, 1, \dots, J-1\}$ with $J < K$. The Fleming-Viot particle system consists of N particles $(\xi_t^v(i))_{i=1}^N$ evolving in \mathcal{A}^c , with a resetting mechanism: particles hitting \mathcal{A} jump to positions of randomly selected surviving particles. This particle system approximately estimates the quasi-stationary distribution

$$v_Q^\pi(x) = \lim_{t \rightarrow \infty} \mathbb{P}^{\partial \mathcal{A}^c}(X_t^\pi = x | T_{\mathcal{K}} > t),$$

where $T_{\mathcal{K}}$ is the hitting time of \mathcal{A} . The FVRL algorithm leverages this to construct gradient estimators:

$$\nabla_{\theta} \widehat{v}^{\pi_{\theta}} = \sum_{x \in \mathcal{A}^c} \widehat{p}^{\pi_{\theta}}(x) \sum_a \widehat{Q}_{\theta}(x, a) \nabla_{\theta} \pi_{\theta}(a | x), \quad (\text{C.3})$$

where $\widehat{p}^{\pi_{\theta}}$ is estimated via the FV particle system.

For $K = 40$, $\rho = 0.7$:

- Vanilla MC estimates $p^\pi(K) \approx 0$ even with 10^6 samples;
- FV provides accurate estimates with $N \sim 10^3$ particles;
- FVRL converges to K^* whereas MC-based policy gradient fails completely.

The method effectively trades the rare-event probability $p^\pi(K)$ for the substantially larger quasi-stationary probability $v_Q^\pi(K)$, overcoming the exponential sample complexity of vanilla approaches. We do not compare here our results with the estimates constructed in Mastropietro et al. (2025) because those are more complex than the idealized situation described in Subsection 3.4, which can, however, serve as a rule of thumb for practitioners.

Endnotes

¹ The processes we consider fall in their class A category with parameter $\alpha = 2$.

² The notation \asymp was introduced in Remark 4.

³ See Borovkov (2013) for the analogous result for discrete-time random walks.

References

- Asmussen S, Glynn PW (2007) *Stochastic Simulation: Algorithms and Analysis*, Stochastic Modelling and Applied Probability, vol. 57 (Springer, New York).
- Asselah A, Ferrari PA, Groisman P (2011) Quasistationary distributions and Fleming-Viot processes in finite spaces. *J. Appl. Probab.* 48(2):322–332.
- Avrachenkov K, Piunovskiy A, Zhang Y (2018) Hitting times in Markov chains with restart and their application to network centrality. *Methodology Comput. Appl. Probab.* 20(4):1173–1188.
- Bertoin J (1998) *Lévy Processes*, Cambridge Tracts in Mathematics, vol. 121 (Cambridge University Press, Cambridge, UK).
- Bertoin J, Doney RA (1994) Cramér’s estimate for Lévy processes. *Statist. Probab. Lett.* 21(5):363–365.
- Borovkov A (2013) *Probability Theory* (Springer, London).
- Budhiraja A, Fraiman N, Waterbury A (2022) Approximating quasi-stationary distributions with interacting reinforced random walks. *ESAIM Probab. Statist.* 26:69–125.
- Burdzy K, Holyst R, Ingeman D, March P (1996) Configurational transition in a Fleming-Viot-type model and probabilistic interpretation of Laplacian eigenfunctions. *J. Phys. A Math. General* 29(11):2633–2642.
- Collet P, Martínez S, San Martín J (2013) *Quasi-stationary Distributions.*, Markov Chains, Diffusions and Dynamical Systems (Springer, Berlin).
- Denisov D, Shneer V (2013) Asymptotics for the first passage times of Lévy processes and random walks. *J. Appl. Probab.* 50(1):64–84.
- Doney RA (1985) Conditional limit theorems for asymptotically stable random walks. *Zeitschrift Wahrscheinlichkeitstheor. Verwandte Gebiete* 70:351–360.
- Doney RA, Maller RA (2004) Moments of passage times for Lévy processes. *Ann. Inst. Henri Poincaré Probab. Statist.* 40(3):279–297.
- Ecoffet A, Huizinga J, Lehman J, Stanley KO, Clune J (2021) First return, then explore. *Nature* 590:580–586.
- Evans MR, Majumdar SN, Schehr G (2020) Stochastic resetting and applications. *J. Phys. A Math. Theoret.* 53(19):193001.
- Ferrari PA, Marić N (2007) Quasi stationary distributions and Fleming-Viot processes in countable spaces. *Electronic J. Probab.* 12:684–702.
- Grigorescu I, Kang M (2013) Markov processes with redistribution. *Markov Processes Related Fields* 19(3):497–520.
- Hansen NR (2009) The maximum of a Lévy process reflected at a general barrier. *Stochastic Processes Appl.* 119(7):2336–2356.
- Höglund T (1990) An asymptotic expression for the probability of ruin within finite time. *Ann. Probab.* 18(1):378–389.
- Iglehart DL (1974) Random walks with negative drift conditioned to stay positive. *J. Appl. Probab.* 11(4):742–751.
- Kolb M, Savov M (2014) Exponential ergodicity of killed Lévy processes in a finite interval. *Electronic Comm. Probab.* 19:1–9.
- Kyprianou AE (2006) *Introductory Lectures on Fluctuations of Lévy Processes with Applications* (Springer, Berlin).
- Kyprianou AE (2014) *Fluctuations of Lévy Processes with Applications*, Introductory Lectures, 2nd ed. (Springer, Berlin).
- Kyprianou AE, Palmowski Z (2006) Quasi-stationary distributions for Lévy processes. *Bernoulli* 12(4):571–581.
- Lambert A (2000) Completely asymmetric Lévy processes confined in a finite interval. *Ann. Inst. Henri Poincaré Probab. Statist.* 36(2):251–274.
- Luby M, Sinclair A, Zuckerman D (1993) Optimal speedup of Las Vegas algorithms. *Inform. Processing Lett.* 47(4):173–180.
- Martínez S, San Martín J (1994) Quasi-stationary distributions for a Brownian motion with drift and associated limit laws. *J. Appl. Probab.* 31(4):911–920.
- Mastropietro D, Ayesta U, Jonckheere M, Majewski S (2025) Fast-exploring reinforcement learning with applications to stochastic networks. *Queueing Systems* 109(3):23.

- Méléard S, Villemonais D (2012) Quasi-stationary distributions and population processes. *Probab. Surveys* 9:340–410.
- Monthus C (2021) Large deviations for Markov processes with stochastic resetting: Analysis via the empirical density and flows or via excursions between resets. *J. Statist. Mech. Theory Experiment* 2021(3):033201.
- Palmowski Z, Pistorius M (2009) Cramér asymptotics for finite time first passage probabilities of general Lévy processes. *Statist. Probab. Lett.* 79(16):1752–1758.
- Villemonais D (2011) Interacting particle systems and Yaglom limit approximation of diffusions with unbounded drift. *Electronic J. Probab.* 16:1663–1692.