Intermittent search, not strict Lévy flight, evolves under relaxed foraging distribution constraints.

Winston Campeau (0000-0001-6032-0186)[1,*]

Andrew M. Simons (0000-0002-0198-465X)[1]

Brett Stevens (0000-0003-4336-1773)[2]

1. Department of Biology, Carleton University, Ottawa, Ontario K1S 5B6, Canada.

2. School of Mathematics and Statistics, Carleton University, Ottawa, Ontario K1S 5B6, Canada.

∗ Corresponding author; e-mail: winston.campeau@carleton.ca.

Manuscript elements: Figure 1, figure 2, figure 3, figure 4, figure 5, figure 6, figure S1, figure S2, figure S3, figure S4, figure S5, supplemental PDF, video S1, video S2, video S3.

Keywords: evolutionary algorithm, optimal foraging, Lévy flight, intermittent search, random walks, search efficiency.

Manuscript type: Article.

Abstract

The survival of an animal depends on its success as a forager, and understanding the adaptations that result in successful foraging strategies is an enduring endeavour of behaviour ecology. Random walks are one of the primary mathematical descriptions of foraging behaviour. Power-law distributions are often used to model random walks as they can characterize a wide range of behaviours, including Lévy walks. Empirical evidence indicates the prevalence and efficiency of Lévy walks as a foraging strategy, and theoretical work suggests an evolutionary origin. However, previous evolutionary models have assumed *a priori* that move lengths are drawn from a power-law or other families of distributions. Here, we remove this restriction with a model that allows for the evolution of any distribution. Instead of Lévy walks, our model unfailingly results in the evolution of intermittent search; a random walk composed of two disjoint modes: frequent localized walks and infrequent extensive moves, which consistently outcompeted Lévy walks. We also demonstrate that foraging using intermittent search may resemble Lévy walk due to interactions with the resources within an environment. These extrinsically-generated Lévy-like walks belie an underlying behaviour and may explain the prevalence of Lévy walks reported in the literature.

Introduction

An important component of an individual animal's fitness depends on the net energy gained through foraging, and the differential fitness of individual animals can depend on the variation among their foraging strategies; more efficient strategies result in more net energy gained thereby increasing the probability of survival (Krebs et al. 1977; Meire and Ervynck 1986; Werner and Hall 1974). Thus, natural selection could favour those strategies which maximize net energy most efficiently. Random walks are often used to model different strategies of animal movement, with move lengths which vary according to an underlying probability distribution. Lévy flight, a type of random walk with move lengths distributed according to the power-law $P(l) \sim l^{-u}$ where $1 < u < 3$, characterized by a mixture of frequent short steps and infrequent longer steps, has been argued to be a particularly efficient searching strategy (Campeau et al. 2022; Viswanathan et al. 1999; Wosniack et al. 2017), and is a searching behaviour that has been observed among many organisms, across multiple taxa (Ariel et al. 2015; Harris et al. 2012; Humphries et al. 2010, 2012; Kölzsch et al. 2015; Lihoreau et al. 2016; Raichlen et al. 2014; Reijers et al. 2019; Reynolds et al. 2018; Shirakawa et al. 2019; Sims et al. 2019, 2014, 2008). Specifically, the Lévy flight foraging hypothesis states that if the searcher has no memory, and resources are scarce, re-visitable, and information of their distribution is unknown to the forager, then a Lévy flight with exponent $u \simeq 2$ is an optimal or near-optimal searching strategy (Viswanathan et al. 2008).

The efficiency and adaptive potential of Lévy flight foraging has been substantiated by both empirical (Humphries et al. 2010, 2012; Lihoreau et al. 2016; Sims et al. 2019, 2008) and theoretical (Campeau et al. 2022; Dannemann et al. 2018; Guinard and Korman 2021; Levernier et al. 2021; Wosniack et al. 2017) evidence; thus, selection could favour Lévy-like behaviour. However, there is an ongoing debate over the relative importance of Lévy-like behaviour as the result of selection (the adaptationist or intrinsic hypothesis) or as an emergent property due to the features of the environment (the emergentist or extrinsic hypothesis). For example, Spider monkeys (*Ate-*

les geoffroyi) exhibit emergent Lévy-like behaviour in response to naturally-occurring Lévy-like tree-size distributions, utilizing mental maps, i.e. memory and perception, when determining which fruiting trees to visit (Boyer et al. 2006). In contrast, free-moving *Drosophila* larvae whose sensory neurons, supraoesophageal ganglion, and brain synaptic activity were blocked exhibited intrinsically generated Lévy-like behaviour (Sims et al. 2019). There is also evidence that some Lévy walks may simply be the by-product of crawling (such as the larvae from Sims et al. (2019)), and are in fact Weierstrassian Lévy walks or stick-slip locomotion (Reynolds 2018, 2021). Whereas there is support for both the extrinsic and intrinsic hypotheses, or even mathematical proof for the optimality of Lévy walks, e.g., Guinard and Korman (2021), the intrinsic hypothesis suggests that natural selection favours Lévy-like behaviour, and poses the additional challenge of requiring evidence of ecological contexts driving its evolution through selection.

There are three notable lines of theoretical evidence which support an evolutionary origin for Lévy flight foraging. Wosniack et al. (2017) demonstrated that among the random walks governed by the the power-law exponents of $1 \leq u \leq 3$, ranging from ballistic ($u = 1$), to Lévy ($u = 2$), to Brownian motion ($u = 3$), a Lévy searching pattern was the most efficient over multiple distinct environments. Dannemann et al. (2018), using the same range of searching patterns, provided evidence that Lévy-like behaviour can maximize the population abundance and minimize the extinctions of a predator-prey system. The combination of Wosniack and Dannemann's research clearly evidence the advantages of Lévy-like behaviour but not evolution through response to selection. One of the first models to demonstrate a response to selection used an evolutionary algorithm where the power-law exponent u, was modeled as a variable and heritable trait (Campeau et al. 2022). The results of the model corroborated both the advantages of Lévy-like behaviour over different environments and its long-term stability, while including both ecological and evolutionary contexts. Although evidence thus far supports the Lévy flight foraging hypothesis, its interpretation is limited by a common constraint; the *a priori* assumption that the searching patterns are governed by a power law. Greater insight into the evolution of for-

aging patterns requires a model that is free of such restrictions on the type of distribution upon which natural selection acts. The *de novo* emergence of Lévy-like distributions would provide a more compelling case of an evolutionary origin.

Here, we use an evolutionary algorithm which includes several important evolutionary ecological contexts from previous models, but dispenses with the *a priori* assumption of search patterns governed by a power-law. In fact, our model makes no assumptions about any family of distributions. Furthermore, our model is meant to be rudimentary in that it isolates the effect of adaptations due to resource encounters, and does not consider interesting extensions such as mate and predator interactions, or a limit to home range. Our model allows for the *de novo* evolution of search patterns with three possible interesting outcomes: 1) search patterns which approximate the resource distributions of the environment, 2) search patterns which approximate a Lévy-like distribution, regardless of the resource distribution and 3) novel search patterns which outcompete at least a Lévy walk. We encode a finite-size list of move lengths as a 'pseudo-genome', where novel genomes are the result of recombination and a round of mutation. The shape of a genome's distribution of move lengths is therefore limited primarily by the size of the genome, thus allowing for a large solution space, and providing firm theoretical grounds for not only testing the Lévy flght foraging hypothesis, but the discovery of any optimal search pattern.

Materials and Methods

We use an evolutionary algorithm to evolve optimal search patterns which may then be compared to existing hypotheses such as the Lévy flight foraging hypothesis. From previous theoretical and mathematical models (Campeau et al. 2022; Dannemann et al. 2018; Wosniack et al. 2017) we include the ecological contexts of population size, lifespan, a proxy for iteroparity (multiple matings over a lifetime), and we vary the environment using different Lévy dust distributions.

We build upon the previous evolutionary models by shifting variation from the level of a distribution parameter (the *a priori* assumption of selection acting on powerlaw exponents; e.g., Campeau et al. 2022; Wosniack et al. 2017), to selection acting on the individual move lengths within a distribution to ask what distribution will result from selection. Evolutionary change in search patterns is driven by 1) the proportion of genetic material in subsequent generations determined by parent individuals' relative fitness, 2) a round of recombination, and 3) a round of mutation. Once a simulation has reached a point of equilibrium, we include an additional set of simulations to assay the fitness of the derived distributions under competition with ancestral and Lévy distributions, and we attempt to parameterize the derived distributions. We also discuss several modifications we have implemented to control for possible biases due to movement occurring on a lattice in the supplementary information.

Environments

A single environment is simulated as an $n \times n$ matrix, or lattice E, with resource entries $e_{i,j}$, and with periodic boundary conditions (i.e., the matrix is mapped onto a torus). Each environment is populated with a sequence of re-visitable resources according to a Lévy dust (LD) distribution; a fractal distribution of resources. The first resource in the sequence is placed at random, with each subsequent resource placed at a distance of l away from the previous resource, where l is randomly selected from the probability mass function (pmf):

$$P(l) = \begin{cases} 0 & \text{if } l < 1 \text{ or } l > n/2 \\ l^{-u} & \text{if } 1 \leq l \leq n/2 \end{cases} \tag{1}$$

We chose $n/2$ as the point of truncation because $n/2$ is the maximum distance between any two points to travel in cardinal directions on a lattice with periodic boundary conditions. The direction of each resource placement in the sequence is selected at random from the set of all locations exactly l distance away. When a resource is placed, the value at the resultant location $e_{i,j}$, is incremented by 1. This process continues until the environment is populated with a sequence of

$n^2 \cdot 10^{-3}$ resources, rounded down to the nearest integer. We chose random exponents $1 \leq u \leq 3$, as well as discrete values $u = 1, 2,$ and 3 for our environments as they represent the distinct distributions produced from ballistic, Lévy, and Brownian movement, respectively (fig. 1).

To increase the environmental variation individuals encounter, and to approximate larger environments, we also utilized a type of procedural generation. A library of environments is produced for each generation of individuals, where each environment has the same amount of resources. Our simulations included libraries where the LD exponent is either: 1) fixed at $u = 1, 2,$ or 3, 2) randomly varied between generations, and 3) randomly varied within each generation. An initial environment is selected at random for each individual. Then, upon crossing the border of an environment, instead of appearing on the opposite side by periodic boundary conditions, a new environment is selected at random from the library and appended to the border. See supplementary information for more details on the procedurally-generated environments.

Evolutionary Algorithm

An evolutionary algorithm is a population-based metaheuristic which borrows evolutionary mechanisms to approximate the solution to an optimization problem (Bäck and Schwefel 1993; Sloss and Gustafson 2020). An individual is a candidate solution to the optimization problem, and the quality of that solution is determined by a fitness function. The following sections will describe our definition of an individual (or candidate solution), the encoding of our evolutionary mechanisms (selection, reproduction, recombination, and mutation), and an outline of the main simulations and subsequent simulations of competition.

Definition of an Individual

An individual was encoded as the following list:

$$[G, pi, pj, \xi, \alpha\xi, \lambda, (I, J), \mathcal{D}]$$

Where G contains the move lengths available to the individual, pi & pj store the position of the individual within an environment, ξ records the number of encountered resources, $\alpha\xi$ is the sum of ξ over time, λ is the lifespan, (I, J) stores the position of the current environment within a lattice of procedurally-generated environments, and $\mathcal{D}$ is a growing dictionary which maps each (I, J) to an environment stored in the library of environments.

G is the 'pseudo-genome'; a finite and unordered list of integer move lengths which can be initialized with any user-defined values. We propose two initial distributions which are interesting from an evolutionary perspective; random samples from a uniform distribution, and an all-ones distribution (equivalent to Brownian motion). All move lengths are equally probable under a uniform distribution; thus, the amount of variation for selection to act upon is maximized. Whereas the all-ones distribution is initially biased towards smaller move lengths and has an initial variance of zero, moving the minimal possible distance is perhaps the simplest initial searching behaviour, barring no movement whatsoever. An individual's initial position (pi, pj) is random, and is updated after every move. Individuals use a truncated random walk; each move is randomly selected from G, and movement occurs until either a resource is located, or the full length of the move has been travelled. Movement occurs on a lattice, and the direction of each move length l is determined by randomly selecting an endpoint from the border of all locations exactly l distance away. An individual will move to the endpoint by taking random selections without replacement from a list of vertical and horizontal increments, continuing until the list is empty. See supplementary information for additional details on the turning algorithm.

The range of perception of an individual is limited to their current cell, or location (i, j). In other words, the "resolution" of the simulation matches the perceptual capabilities of an individual, and was chosen to intentionally minimize their information on the distribution of resources.

When an individual encounters a resource at a location (i, j), ξ is incremented by $e_{i,j}$ (a growing sum), whereas $\alpha\xi$ is incremented by ξ after every entry-wise move, regardless of resource encounter. Thus, ξ measures the total number of resources found, and $\alpha\xi$ is used to measure the average number of resources found over the lifespan of an individual. The lifespan of an individual λ, is the total distance (or total time) available until that individual is removed from the simulation. We normalized lifespans by the size of the environment according to $\varsigma \cdot n^2$, i.e., each individual can theoretically visit every location of a single environment, ς times (by default $\varsigma = 20$). To track an individual's location within their procedurally-generated environments, (I, J) is the index of their current environment, which is updated when crossing a boundary of their current environment, and is stored as a key in a growing dictionary $\mathcal{D}$. Each key stored in dictionary maps to an environment in the library of environments, and preserves the relative locations of the procedurally-generated environment for that individual. See supplementary information for more details on the procedural-generation of environments.

Selection or Fitness Function

Once all individuals have exhausted their lifespans, their resultant average number of resources found over their lifespans ($\alpha\xi$) are consigned to the fitness function. There is also the option to use the total number of resources located (ξ) as a measure of fitness, but we used $\alpha\xi$ for two reasons. The first is that ξ and $\alpha\xi$ are strongly correlated; thus, fitness differences are minimal (Campeau et al. 2022). The second reason is that the average energy over a lifespan $\alpha\xi$, would be markedly different for an individual that locates all of its resources at the beginning of its lifespan rather than at the end, but this would be indistinguishable using ξ, the energy at the end of a lifespan. Thus, ξ resembles semelparity because of a lifetime of energy invested into reproduction only at the end of a lifespan, whereas $\alpha\xi$ resembles a proxy for iteroparity because it considers the average amount of energy available over a lifespan, and could be interpreted as the energy available for multiple rounds of reproduction. Organisms also more frequently evolve iteroparous reproduction strategies (Hughes 2017), and many semelparous strategies ap-

proximate iteroparity over shorter timescales (Hughes and Simons 2014), thus $\alpha\xi$ is arguably the more appropriate choice.

The fitness function uses the $\alpha\xi$ values to determine the number of gametes that each parent will contribute to form the subsequent generation according to:

$$\omega_i = \left\lceil \frac{K \cdot \alpha\xi_i}{\sum_{j=0}^{\psi} \alpha\xi_j} \right\rceil \tag{2}$$

where for each parent, each half increment of ω_i represents one gamete, and the total number of gametes is based on the ratio of the carrying capacity (K) and their individual fitness relative to the sum of all surviving individuals (rounded up to the nearest integer). We index parents until ψ, instead of K, because an individual does not "survive" if its $\alpha\xi \leq 0$; thus, ψ is the number of surviving individuals. Each parent is assigned a value of ω_i, in descending order of $\alpha\xi$, until there are enough gametes to generate the next generation (i.e., $\sum \omega_i = K$). Therefore, although an individual may have the potential to have $\omega_i > 0$, there remains the possibility that its genetic information is not included in the reproductive pool. The resultant output is a list of pairs $(G_i, \omega_i) \in \Omega$, which are used to generate the next generation of individuals.

Reproduction Through Recombination and Mutation

Mating is random, syngamous, and without incompatible mating types. Two parent pairs (G_i, ω_i) & (G_j, ω_j) are randomly selected from Ω, and upon selection, each parental genome (G_i & G_j) is first randomly rearranged (i.e., no linkage). One half of each parental genome is then assigned at random as its gametic contribution to the single offspring, G_k (fig. 2). Accordingly, ω_i and ω_j are then decremented by one half. Whereas we recognize the differential energetic burdens of reproduction and its evolutionary importance, we believe its inclusion would add unnecessary complexity in the scope of this paper.

The resultant offspring genome then undergoes a round of mutation. First, a fixed num-

ber of move lengths are randomly selected for mutation. Each move length is then modified by a random value m, from a scaled and discretized unit normal distribution. We discretize by rounding to comply with the integer requirement, and use a scalar η, to vary the magnitude of mutation. We extend the same bound restrictions from equation (1), such that a mutated move length cannot go below one, and cannot exceed $n/2$. Mutations resulting in move lengths below one or above $n/2$ will roll back into the space of $1 \leq l \leq n/2$, where the amount of roll back is the difference over the bound. For example, if $n/2 = 100$, then $l = 110$ would become $l = 90$, and $l = -10$ would become $l = 11$. The value of η should be carefully selected as to not result in mutations so large that the resultant move lengths bounce over the bounds until reaching an equilibrium ($5\eta < n/2$ is a sufficient condition to avoid this behaviour). The process of reproduction, recombination, and mutation continues until all $\omega_i = 0$. In the case where there is only one parent remaining, reproduction switches to asexual, and the offspring genome is a mutated copy of the parent genome.

The simulation is initialized by choosing an exponent(s) for the distribution of resources, the maximum number of generations, the mutation rate and magnitude, the size of the environment and the library which contains them, and the size of the population and their genomes. The library of environments is generated first, with each environment receiving $n^2 \cdot 10^{-3}$ resources, rounded down to the nearest integer. Next, the starting population is generated, with a specified initial distribution for G; all-ones or random uniform move lengths for this study. The individuals traverse procedurally-generated environments until all of their lifespans are exhausted. The population is then assessed by the fitness function, and the offspring are generated by the recombination and mutation of random selections from any two parents who were admitted to the reproductive pool. This process continues until the number of offspring is equal to the selected population size. The offspring become the subsequent population, and the process is repeated until either the mean population $\alpha\xi$ stabilizes or the total number of generations is exhausted (10^4 by default).

All genomes are saved every g^{th} generation (by default, $g = 1$) as a "fossil record", similar to the fossil record in Lenski's Long-Term Evolution Experiment (Wiser and Lenski 2015), allowing the simulation to be restarted at any point, if necessary. This data also allows for competitive fitness assays between our ancestral and derived foraging patterns, and further, assays of the derived foraging patterns competing with Lévy foraging patterns to determine quantitative differences, and provide qualitative insight. For the competition simulations, the initial population was composed of either derived and ancestral individuals, or derived and Lévy individuals, in equal frequency. In these competition assays, reproduction was strictly asexual; parents produced offspring with exact copies of their genomes. And lastly, the two populations competed for a fraction of a carrying capacity, and a simulation was stopped once an equilibrium was reached (i.e., only one type remained).

Simulations and results were programmed using Python 3.10.6 and the numpy (Harris et al. 2020), scipy (Virtanen et al. 2020), pandas (Wes McKinney 2010), multiprocessing (McKerns et al. 2012), and Matplotlib (Caswell et al. 2021) libraries in addition to the Python standard library. Computations were performed on the Graham, Cedar, Narval, and Niagara supercomputers at the Digital Research Alliance of Canada (Loken et al. 2010; Ponce et al. 2019).

Distribution Characterization

Our evolutionary algorithm constrains the evolution of distributions to containing s move lengths (the total number of move lengths in G) where each move length cannot exceed a value of $n/2$. This means our algorithm can result in a potential

$$\binom{\frac{n}{2} + s - 1}{s}$$

unique distributions. For example, $s = 100$ and $n/2 = 100$ is approximately $4.5 \cdot 10^{58}$ possibilities and, albeit not without many correlated distributions, it is difficult to anticipate exactly what

methods would be necessary to characterize the derived distributions (i.e., there is no guarantee whether distribution fitting will be informative). Based on visual inspection of the derived distributions and their behaviour, we considered comparisons with an array of candidate models of random walks (e.g., composite random walks or power-laws), and measured their scaling exponents by their mean-squared displacement.

Results

We ran simulations where populations of 100 individuals would search for re-visitable resources within a procedurally-generated environment. Their initial environment was encoded as a matrix $E_{n\times n}$, with dimension $n = 997$ (see supplementary information for why n is prime). Each matrix had resources distributed by either successive ballistic, Lévy, or Brownian flights (fig. 1), or mixtures of resource distributions with random LD exponents. Simulating evolution over these distinct distributions enabled us to test whether the derived distributions would converge to environmental, Lévy-like, or novel distributions. Each individual was assigned a lifespan of $\lambda = 20 \cdot n^2$ steps, and an initial G composed of either 200 random uniform move lengths $1 \leq l \leq n/2$ or 200 ones. We ran simulations with each initial G for the environments with discrete LD exponents, and with only the random uniform G for environments with randomly varied LD exponents between and within generations. We ran replications for only the random uniform starting G, thus resulting in a total of 13 evolutionary simulations. We set a fixed mutation rate such that 12 out of 200 move lengths were always admitted for mutation, with magnitude $\eta = 2$. Once the evolution of G stabilized, we then competed the derived distributions against ancestral distributions to test for adaptive evolution, and against Lévy distributions to determine quantitative and qualitative differences, if any.

Distribution Evolution & Characterization

All simulations resulted in similar derived distributions, and regardless of their environment (fig. 3, video S1, video S2). The random uniform and all-ones starting conditions required approximately 400 and 5000 generations to reach equilibration, respectively. The derived distributions are bimodal with the primary mode clustered near one, and a secondary mode clustered on the $n/2$ point of truncation. The resultant behaviour of the bimodal distributions is an emergent two-phase walk (fig 4CDE, video S3A). The primary phase is highly localized, and visually similar to a Brownian walk. The secondary phase could be described as a ballistic or scaled Brownian walk, where the most common move length was $l \approx n/2$ constituting $\approx 17\%$ of 18 unique values, although the size of the secondary modes (2.5 move lengths on average) were too small for distribution fitting. From the non-continuity of the derived distributions and visual inspection, we proposed two candidate composite random walks:

$$p(l) = p_{primary}((1 - p_{geometric})^{l-1} p_{geometric}) + (1 - p_{primary})(n/2) \tag{3}$$

$$p(l) = p_{primary}(l^{-u}) + (1 - p_{primary})(n/2), \quad 1 \leq u \leq 5 \tag{4}$$

where $p_{primary}$ is the probability of drawing a move length from the primary mode, modelled by either a geometric distribution (the discrete equivalent of an exponential distribution) with $p_{geometric}$, or a power law with $1 \leq u \leq 5$, and a $1 - p_{primary}$ probability of drawing from the secondary mode which is simply a move length of $l = n/2$. To determine the suitability of each model we first isolated the primary modes from each of our 1300 derived distributions and compared their cumulative distribution functions (CDF) over an array of 10^3 power-law exponents u from $[1.0, 5.0]$ and $p_{geometric}$ from $[0, 1]$ using k-sample Anderson-Darling (AD) tests. The null hypothesis of the k-sample AD test is that two (or more) samples are drawn from the same distribution; a power law with $u \approx 1.9$ resulted in a failure to reject the null hypothesis for $\approx 1\%$ of the comparisons, whereas a geometric distribution with $p_{geometric} \approx 0.33$ resulted in a failure to

reject the null hypothesis for $\approx 58\%$ of the comparisons. Randomly adding either two or three move lengths of $l = n/2$ to the optimal geometric distribution resulted in a failure to reject the null for $\approx 60\%$ of the comparisons with the complete derived distributions. The behaviour of the resultant model was then visually compared with a distribution derived from each environment (fig. 4, video S3AB).

Next, we computed the scaling exponent of one of the derived populations, and included a Lévy walk (as defined by the pmf in equation (1) for comparison. The scaling exponent is obtained by measuring the displacement of multiple individuals over increasing periods of time, taking the mean of their squared displacements, and observing how those mean-squared displacements (MSD) change over time, i.e., measuring the slope. The scaling exponent can be used to compare and characterize the exploratory behaviour of different random walks (Einstein 1905; Frenkel and Smit 2023). A simple Brownian walk has a scaling exponent of $\alpha = 1$, indicating a walk which displaces linearly in time, whereas the scaling exponents for the Lévy walks and derived walks, computed using a least-squares Levenberg-Marquardt algorithm based on the statistical regularity of the data and its use for similar analysis (Liu et al. 2021; Zhou et al. 2018), were found to be $\alpha = 1.97$ ($R^2 = 0.999$) and $\alpha = 1.997$ ($R^2 = 0.999$), respectively (fig. S1). These values indicate that individuals using either strategy will explore the environment much faster than a Brownian walk, and that the derived walks will do so marginally faster than Lévy walks, but are otherwise difficult to distinguish by their scaling exponents, alone.

Lastly, we performed an exploratory analysis to examine the 'realized' distributions of move lengths. We define a realized distribution as the result of the truncated move lengths, rather than the lengths pulled from a distribution. For example, an individual might pull $l = 498$ from their distribution but encounter a resource after only 10 steps; thus, the realized length is the truncated length of 10 steps. We collected the realized distributions from each type of environment, over a range of resource densities (the tested $n^2 \cdot 10^{-3}$, and one fold up & down) and a lifespan

of $\lambda = n^2$ (the minimum to span the entire environment) to determine how they might differ from the derived distributions. The distributions of realized move lengths were increasingly contiguous with increasing resource density, and decreasing resource clumpiness, i.e., $u < 3$ (fig. 5). We also measured the degree of contiguousness by sorting the realized distributions in increasing order and measuring the largest gap (any difference greater than one between successive move lengths). Whereas the largest gap of the derived distributions tended towards $(n/2) - 1$, the realized distributions were either contiguous or had gaps no larger than size four. Assuming a power-law fit almost always resulted in exponents falling in the "Lévy-like" range, with $1 < u < 3$, often close to a "true" Lévy walk of $u = 2.0$ (fig. S2).

Competition Simulations

We computed a sum total of 324 competition simulations between the derived distributions and either a Lévy walk with $u = 2.0$ or root and mid-generation ancestors from the simulated fossil records. Competition between derived and Lévy occurred over environments $LD = 1.0$, 2.0, and 3.0 with either $n^2 \cdot 10^{-3}$ or $n^2 \cdot 10^{-4}$ resources, with lifespans of $\lambda = 20 \cdot n^2$ or $\lambda = n^2$, starting with either populations of 10 or 100 each (i.e., competing for carrying capacities of either 20 or 200), and with the Lévy walks having access to a G of size 200 or 2000. Three replicates were computed for each of these 48 possible combinations, and therefore totalling 144 simulations for a single population of derived distributions. The derived populations from our six evolutionary simulations were near indistinguishable, so we replicated the competition simulations using only two of the derived populations: individuals evolved from all-ones distributions and LD $u = 2.0$ and random uniform distributions and LD $u = 1.0$, thus totalling 288 simulations. We competed the derived distributions against their root ancestor (i.e., the 0^{th} generation), and a mid-generation distribution (100^{th} for the random uniform case and 2000^{th} for all ones), under the same conditions in which they evolved, and with three replicates each, totalling 36 simulations.

The derived populations won 100% of the 324 competition simulations. The median number of generations required to displace the Lévy populations across all simulations was two. The largest source of variance in generational time among the simulations were due to lifespan, and the derived population (fig. 6). Lifespans of $\lambda = 20 \cdot n^2$ required a median of two generations (min=1, max=3), whereas $\lambda = n^2$ took three generations (min=1, max=11). Within the lifespan of $\lambda = 20 \cdot n^2$, the initial all-ones distribution required a median of two generations (min=1, max=3), whereas the initial random uniform distributions required one (min=1, max=2), and within the lifespan of $\lambda = n^2$ the initial all-ones distributions required a median of four generations (min=1, max=11), whereas the initial random uniform distributions required three generations (min=1, max=6). We also include a subset of simulations using very small lifespans $\lambda = 0.1 \cdot n^2$, available in the supplementary information. The overall median number of generations required to displace the ancestral populations was two (min=1, max=2). Although the derived distributions were largely indistinguishable, the differences in generational time among the two starting conditions to displace the Lévy populations motivated an additional set of analyses and simulation to explain those differences, and these results are available in the supplementary information.

In a final, and additional, set of simulations we assayed the success of the derived populations against Lévy populations scaled with n. We scaled the environment up and down by a factor of two ($2n = 1994$ and $n/2 = 498$), and set the: lifespan $\lambda = 20 \cdot n^2$, patch density $10^{-3} \cdot n^2$, and a G of size 2000 for the Lévy individuals, and over each environment type, with three replications. We adjusted the derived distributions by simply scaling only the secondary mode by the same factor of two. The scaled derived distributions won 100% of these 18 additional simulations. The median number of generations required to displace the Lévy distributions in the upscaled environments was two (min=2, max=2), and a median of three generations for the downscaled environments (min=2, max=3).

Discussion

The Lévy flight foraging hypothesis, in its simplest form, is the claim that because Lévy flight is an optimal or at least adaptive searching strategy, then natural selection should result in Lévy-like behvaiour. Empirical evidence has indicated that many organisms exhibit Lévy-like behaviour, but whether that behaviour is intrinsic and the result of natural selection, is extrinsic and an emergent phenomenon due to interactions with the environment, or even a mixture of both, is unclear. Lévy walks have been evidenced as adaptive over various environments (Wosniack et al. 2017), a strategy which increases the stability of Lotka-Volterra systems (Dannemann et al. 2018), and as a response to selection regardless of the fractal dimension of resources and under several evolutionary and ecological contexts (Campeau et al. 2022). However a common limitation of these studies is the *a priori* assumption that foragers adhere to search patterns governed by power-laws. Building upon existing models, we removed this limitation and constructed an evolutionary algorithm which allowed for the *de novo* evolution of search patterns; we shifted the resolution on which evolutionary mechanisms could act on from the parameters of a distribution to the distributions themselves. We proposed three possible interesting outcomes: 1) the patterns approximate their environment, 2) the patterns converge to Lévy-like distributions, or 3) novel search patterns which also outcompete at least a Lévy walk. Our results unequivocally point to the third outcome.

The search patterns of our derived populations, regardless of their initial conditions, all converged to a search comprised of two emergent phases: a long & highly localized phase and a shorter ballistic phase, where under competition they outcompeted Lévy walks, including in scaled environments. This search pattern is not new to the literature, and is commonly referred to as intermittent search (Bénichou et al. 2005, 2011; Bénichou et al. 2006), a composite random walk which is known to outcompete Lévy-like behaviour (Benhamou and Collet 2015), and characterizes the behaviour of some organisms (Reynolds and Frye 2007). The advantage of intermittent

search should feel intuitive; when searching for food (e.g., mushrooms) or something lost, humans will sometimes search a small area intensively before eventually giving up and switching to a new location (Pacheco-Cobos et al. 2019). The primary mode of the derived distributions, composed of smaller moves, provides the capability to search intensively and exploit clustered resources. The secondary mode, composed of a few large moves, allows an individual to quickly explore the environment before eventually encountering another resource. We also conjecture that limiting the size of the larger moves may be unnecessary. For example, it may be sufficient to simply continue walking until a resource is located, which does not require a predetermined distance, but accomplishes the same result.

A similar search, area-restricted search (Dorfman et al. 2022; Kareiva and Odell 1987), can be characterized by an intermittent search as foragers switch between an intense localized search and an exploratory global phase. The primary difference between the two strategies being that turning angles increase following resource encounters under an area-restricted search – our evolved walks do not react in response to resource encounters, they are simply random draws from a single distribution. Evolutionary algorithms and models of reinforcement learning have previously been used to evidence the adaptive value of area-restricted search, but these models either acted primarily on the duration spent in the local search, and the directionality of the local and global phases (therefore presupposing the option of intermittent search; Hills 2006; Scharf et al. 2009), or included an explicit capability to learn (López-Incera et al. 2020). In contrast, our derived distributions are the result of combinatorial optimization from random processes in the absence of sensory or memory effects; each distribution G is a multiset of size 200, with move lengths belonging to a finite set of size 498 (the restriction set by equation (1), thus there were a possible $697!/(200!497!)$ unique distributions – more than the sum of both the legal and theoretical number of positions is the game 'Go' (Tromp 2016; Walraet and Tromp 2016). Further, we utilized a form of procedurally-generated environments, whose increased environmental variation decreased the risk of artifacts that might incur from a single environment with periodic

boundary conditions. Thus, the consistent result of intermittent search in all of our simulations is a strong indication of the selective advantage of intermittent search as an intrinsically-generated behaviour.

If intermittent search is intrinsically-generated, then its behaviour may only ever be observed within a vacuum, and may be mistaken for Lévy-like behaviour, since interactions with the environment will inevitably affect the observed searching behaviour. For example, 'brain-blocked' free-moving *Drosophila* larvae exhibited movement similar to a Lévy walk with an power-law exponent of 1.96 (Sims et al. 2019). By eliminating the larvae's sensory information of the environment, an alluring hypothesis is that their movement was an intrinsically-generated Lévy walk. However, the larvae could collide with one another and the edges of their arena, and the agar coating the arena may have not been featureless. We demonstrated this possibility by allowing the derived individuals to search within environments with various distributions and densities of resources and collecting the 'realized' move lengths (fig. 5). Whereas the evolved distributions were bimodal and disjoint, the observed distributions were contiguous, heavy-tailed, and under the assumption of a power law exhibited exponents which fell within the range of "Lévy-like" behaviour, $1 < u < 3$ (fig. S2). The MSD scaling exponent of the derived behaviour, $\alpha = 1.997$, was nearly identical to that of a Lévy-walk, $\alpha = 1.97$, and scaling exponents have been reported in the literature as evidence to characterize organisms as using Lévy walk, with many $1.6 \leq \alpha \leq 1.9$ (Ariel et al. 2015; Miramontes et al. 2014; Murakami et al. 2015; Ramos-Fernandez et al. 2003), further conflating the two behaviours. The scaling exponent of our derived individuals is also a near exact match with the results of the model from López-Incera et al. (2020), despite our unique assumptions.

Although our results strongly indicate the adaptiveness of intermittent search, we recognize that our model is chasing after an optimal and potentially universal strategy that may not exist. Selection acts on preexisting variation, and intermittent search may be unattainable or costly

to evolve in real-life organisms, despite its advantages. Thus, we do not necessarily expect intrinsically-generated intermittent search, or any other distribution generated by our model, to be a universally-available strategy. However, with the appropriate ecological contexts, and given the evolutionary nature of the model's heuristics, we believe our model has the potential to reveal a class of biologically-relevant and intrinsically-generated search patterns, and that includes intermittent search.

Conclusion

We overcame the *a priori* constraint of searching behaviour governed by power laws, and constructed a model which allowed for the evolution of virtually any probability distribution. Whereas the Lévy flight foraging hypothesis predicts the evolution of adaptations for Lévy-like behaviour, our model instead resulted in intermittent search, a behaviour which characterizes an area-restricted search. Our results do not eliminate the possibility of intrinsically-generated Lévy-like behaviour, but they do evidence the optimality of an intrinsically-generated intermittent searching behaviour – a behaviour which outcompeted Lévy walks in all 324 of our tested instances, most often only requiring two generations. We also demonstrated that intermittent search could be perceived as Lévy-like behaviour due to interactions with the environment, thus supporting aspects of the extrinsic hypothesis. The results of our model could be confirmed with experiments which minimize or eliminate truncated movement due to interactions with the environment. Further, if intermittent search is intrinsically-generated, then the Fourier analysis of neuronal activity in, say, peristaltic-driven motion (e.g., pedal waves in larvae as described by Reynolds (2021)), might reveal two primary oscillatory components: a low-frequency & large-amplitude component, and a high-frequency & low-amplitude component which accounts for the majority of the power spectrum. Our model, as presented here, is most appropriate for exploring the evolution of intrinsically-generated behaviour, specifically in the absence of information on the environment. The model could be extended by subjecting the distribution of turning angles to

selection, or including aspects of memory (working memory, as opposed to a genetic memory) to explore more complex behaviours. For example, turning angles are not independent of resource encounters under area-restricted search (Dorfman et al. 2022), and selection could act on the correlation of turning angles and resource encounters. The evolutionary mechanisms could include saltatory mutations (i.e., mutation probabilities other than normal), and considerations for the relatedness of individuals to alter how the algorithm explores the solution space. The model could also include predator and mate encounters to determine how combinations of encounters affect the evolution of individual walks. The environments could include spatiotemporal fluctuations in resource density to assess how food availability alters searching behaviour. And lastly, our model could include costs for cognitive or physiological adaptations, and explore how those costs might affect the evolution of intermittent search or whichever resultant evolved walks. We hope the results presented here will encourage further research on the potential evolutionary origins of an intrinsic and intermittent search.

Acknowledgements

Thank you to R. Taylor for providing additional computational resources, to F. Lutscher and R. Gorelick for their insight and advice, to A. Cheslock, I. Berberi, and M. Coffey for their support, and to the editors and referees for their detailed and helpful suggestions. This work was supported by the Natural Sciences and Engineering Research Council of Canada: RGPIN-2021-03832 (AMS) and RGPIN-2017-06392 (BS).

Statement of Authorship

All authors contributed to the original conception of the model and to the development of the experimental methods. WC wrote the original draft, coded the simulations, and collected, analyzed, validated and prepared visualizations of the resultant data. All authors reviewed and edited the writing at all stages of revision.

Data and Code Availability

All data and code have been deposited in the Zenodo data repository (https://zenodo.org/doi/10.5281/zenodo.7971749; Campeau et al. 2023).

Literature Cited

Ariel, G., A. Rabani, S. Benisty, J. D. Partridge, R. M. Harshey, and A. Be'er. 2015. Swarming bacteria migrate by Lévy walk. Nature Communications 6:8396. Number: 1 publisher: Nature Publishing Group.

Bäck, T., and H.-P. Schwefel. 1993. An overview of evolutionary algorithms for parameter optimization. Evol. Comput. 1:1–23.

Benhamou, S., and J. Collet. 2015. Ultimate failure of the Lévy foraging hypothesis: Two-scale searching strategies outperform scale-free ones even when prey are scarce and cryptic. Journal of Theoretical Biology 387:221–227.

Bénichou, O., M. Coppey, M. Moreau, P.-H. Suet, and R. Voituriez. 2005. Optimal search strategies for hidden targets. Phys. Rev. Lett. 94:198101.

Bénichou, O., C. Loverdo, M. Moreau, and R. Voituriez. 2011. Intermittent search strategies. Rev. Mod. Phys. 83:81–129.

Boyer, D., G. Ramos-Fernandez, O. Miramontes, J. Mateos, C. Germinal, H. Larralde, H. Ramos, and F. Rojas. 2006. Scale-free foraging by primates emerges from their interaction with a complex environment. Proceedings. Biological sciences / The Royal Society 273:1743–50.

Bénichou, O., M. Coppey, M. Moreau, and R. Voituriez. 2006. Intermittent search strategies: When losing time becomes efficient. Europhysics Letters 75:349.

Campeau, W., A. Simons, and B. Stevens. 2023. Data For: Intermittent search, not strict Lévy flight, evolves under relaxed foraging distribution constraints.

Campeau, W., A. M. Simons, and B. Stevens. 2022. The evolutionary maintenance of Lévy flight foraging. PLOS Computational Biology 18:1–26.

Caswell, T. A., M. Droettboom, A. Lee, d. E. S. Andrade, J. Hunter, T. Hoffmann, E. Firing, J. Klymak, D. Stansby, N. Varoquaux, J. H. Nielsen, B. Root, R. May, P. Elson, J. K. Seppänen, D. Dale, J.-J. Lee, D. McDougall, A. Straw, P. Hobson, C. Gohlke, T. S. Yu, E. Ma, A. F. Vincent, hannah, S. Silvester, C. Moad, N. Kniazev, E. Ernest, and P. Ivanov. 2021. matplotlib/matplotlib: REL: v3.4.0rc3. Zenodo. DOI: 10.5281/zenodo.4595937.

Dannemann, T., D. Boyer, and O. Miramontes. 2018. Lévy flight movements prevent extinctions and maximize population abundances in fragile lotka-volterra systems. Proceedings of the National Academy of Sciences 115:3794–3799.

Dorfman, A., T. T. Hills, and I. Scharf. 2022. A guide to area-restricted search: a foundational foraging behaviour. Biological Reviews 97:2076–2089.

Einstein, A. 1905. About the movement of suspended particles in liquids at rest as required by the molecular kinetic theory of heat.[uber die von der molekularkinetischen theorie der warme geforderte bewegung von in ruhenden flussigkeiten suspendierten teilchen]. Ann Phys 322:549–560.

Frenkel, D., and B. Smit. 2023. Understanding Molecular Simulation: From Algorithms to Applications. Elsevier Science.

Guinard, B., and A. Korman. 2021. Intermittent inverse-square lévy walks are optimal for finding targets of all sizes. Science Advances 7:eabe8211.

Harris, C. R., K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett,

A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant. 2020. Array programming with numpy. Nature 585:357–362. Number: 7825 publisher: Nature Publishing Group.

Harris, T. H., E. J. Banigan, D. A. Christian, C. Konradt, E. D. Tait Wojno, K. Norose, E. H. Wilson, B. John, W. Weninger, A. D. Luster, A. J. Liu, and C. A. Hunter. 2012. Generalized Lévy walks and the role of chemokines in migration of effector $cd8^{+}t$ cells. Nature 486:545–548. Number: 7404 publisher: Nature Publishing Group.

Hills, T. T. 2006. Animal foraging and the evolution of goal-directed cognition. Cognitive Science 30:3–41.

Hughes, P. W. 2017. Between semelparity and iteroparity: Empirical evidence for a continuum of modes of parity. Ecology and Evolution 7:8232–8261.

Hughes, P. W., and A. Simons. 2014. The continuum between semelparity and iteroparity: Plastic expression of parity in response to season length manipulation in lobelia inflata. BMC evolutionary biology 14:90.

Humphries, N. E., N. Queiroz, J. R. M. Dyer, N. G. Pade, M. K. Musyl, K. M. Schaefer, D. W. Fuller, J. M. Brunnschweiler, T. K. Doyle, J. D. R. Houghton, G. C. Hays, C. S. Jones, L. R. Noble, V. J. Wearmouth, E. J. Southall, and D. W. Sims. 2010. Environmental context explains Lévy and Brownian movement patterns of marine predators. Nature 465:1066–1069. Number: 7301 publisher: Nature Publishing Group.

Humphries, N. E., H. Weimerskirch, N. Queiroz, E. J. Southall, and D. W. Sims. 2012. Foraging success of biological Lévy flights recorded in situ. Proceedings of the National Academy of Sciences of the United States of America 109:7169–7174.

Kareiva, P., and G. Odell. 1987. Swarms of predators exhibit ”preytaxis” if individual predators use area-restricted search. The American Naturalist 130:233–270.

Krebs, J. R., J. T. Erichsen, M. I. Webber, and E. L. Charnov. 1977. Optimal prey selection in the great tit (parus major). Animal Behaviour 25:30–38.

Kölzsch, A., A. Alzate, F. Bartumeus, M. de Jager, E. J. Weerman, G. M. Hengeveld, M. Naguib, B. A. Nolet, and J. van de Koppel. 2015. Experimental evidence for inherent Lévy search behaviour in foraging animals. Proceedings of the Royal Society B: Biological Sciences 282:20150424. Publisher: Royal Society.

Levernier, N., O. Bénichou, and R. Voituriez. 2021. Universality classes of hitting probabilities of jump processes. Phys. Rev. Lett. 126:100602.

Lihoreau, M., T. C. Ings, L. Chittka, and A. M. Reynolds. 2016. Signatures of a globally optimal searching strategy in the three-dimensional foraging flights of bumblebees. Scientific Reports 6.

Liu, Y., X. Long, P. Martin, S. Solomon, and P. Gong. 2021. Lévy walk dynamics explain gamma burst patterns in primate cerebral cortex. Communications Biology 4:739.

Loken, C., D. Gruner, L. Groer, R. Peltier, N. Bunn, M. Craig, T. Henriques, J. Dempsey, C.-H. Yu, J. Chen, L. J. Dursi, J. Chong, S. Northrup, J. Pinto, N. Knecht, and R. V. Zon. 2010. SciNet: Lessons learned from building a power-efficient top-20 system and data centre. Journal of Physics: Conference Series 256:012026.

López-Incera, A., K. Ried, T. Müller, and H. J. Briegel. 2020. Development of swarm behavior in artificial learning agents that adapt to different foraging environments. PLOS ONE 15:1–38.

McKerns, M. M., L. Strand, T. Sullivan, A. Fang, and M. A. G. Aivazis. 2012. Building a framework for predictive science. CoRR abs/1202.1056.

Meire, P., and A. Ervynck. 1986. Are oystercatchers (haematopus ostralegus) selecting the most profitable mussels (mytilus edulis)? Animal Behaviour 34:1427–1435.

Miramontes, O., O. DeSouza, L. R. Paiva, A. Marins, and S. Orozco. 2014. Lévy flights and self-similar exploratory behaviour of termite workers: Beyond model fitting. PLOS ONE 9:1–9.

Murakami, H., T. Niizato, T. Tomaru, Y. Nishiyama, and Y.-P. Gunji. 2015. Inherent noise appears as a Lévy walk in fish schools. Scientific reports 5:10605.

Pacheco-Cobos, L., B. Winterhalder, C. Cuatianquiz Lima, M. Rosetti, R. Hudson, and C. Ross. 2019. Nahua mushroom gatherers use area-restricted search strategies that conform to marginal value theorem predictions. Proceedings of the National Academy of Sciences 116:201814476.

Ponce, M., R. van Zon, S. Northrup, D. Gruner, J. Chen, F. Ertinaz, A. Fedoseev, L. Groer, F. Mao, B. C. Mundim, M. Nolta, J. Pinto, M. Saldarriaga, V. Slavnic, E. Spence, C.-H. Yu, and W. R. Peltier. 2019. Deploying a top-100 supercomputer for large parallel workloads: The niagara supercomputer. *In* Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (Learning). PEARC '19. Association for Computing Machinery, New York, NY, USA.

Raichlen, D. A., B. M. Wood, A. D. Gordon, A. Z. P. Mabulla, F. W. Marlowe, and H. Pontzer. 2014. Evidence of Lévy walk foraging patterns in human hunter–gatherers. Proceedings of the National Academy of Sciences of the United States of America 111:728–733.

Ramos-Fernandez, G., J. Mateos, O. Miramontes, C. Germinal, H. Larralde, and B. Ayala-Orozco. 2003. Lévy walk patterns in the foraging movements of spider monkeys (ateles geoffroyi). Behavioral Ecology and Sociobiology 55.

Reijers, V. C., K. Siteur, S. Hoeks, J. van Belzen, A. C. W. Borst, J. H. T. Heusinkveld, L. L. Govers, T. J. Bouma, L. P. M. Lamers, J. van de Koppel, and T. van der Heide. 2019. A Lévy expansion strategy optimizes early dune building by beach grasses. Nature Communications 10:2656. Number: 1 publisher: Nature Publishing Group.

Reynolds, A., E. Ceccon, C. Baldauf, T. Karina Medeiros, and O. Miramontes. 2018. Lévy foraging patterns of rural humans. PLoS ONE 13.

Reynolds, A. M. 2018. Current status and future directions of Lévy walk research. Biology Open 7:bio030106.

———. 2021. Weierstrassian Lévy walks are a by-product of crawling. Eur. Phys. J. E 44:96.

Reynolds, A. M., and M. A. Frye. 2007. Free-flight odor tracking in drosophila is consistent with an optimal intermittent scale-free search. PLOS ONE 2:1–9.

Scharf, I., B. Kotler, and O. Ovadia. 2009. Consequences of food distribution for optimal searching behavior: an evolutionary model. Evolutionary Ecology 23:245–259.

Shirakawa, T., T. Niizato, H. Sato, and R. Ohno. 2019. Biased Lévy-walk pattern in the exploratory behavior of the physarum plasmodium. Biosystems 182:52–58.

Sims, D. W., N. E. Humphries, N. Hu, V. Medan, and J. Berni. 2019. Optimal searching behaviour generated intrinsically by the central pattern generator for locomotion. eLife 8.

Sims, D. W., A. M. Reynolds, N. E. Humphries, E. J. Southall, V. J. Wearmouth, B. Metcalfe, and R. J. Twitchett. 2014. Hierarchical random walks in trace fossils and the origin of optimal search behavior. Proceedings of the National Academy of Sciences of the United States of America 111:11073–11078.

Sims, D. W., E. J. Southall, N. E. Humphries, G. C. Hays, C. J. A. Bradshaw, J. W. Pitchford, A. James, M. Z. Ahmed, A. S. Brierley, M. A. Hindell, D. Morritt, M. K. Musyl, D. Righton, E. L. C. Shepard, V. J. Wearmouth, R. P. Wilson, M. J. Witt, and J. D. Metcalfe. 2008. Scaling laws of marine predator search behaviour. Nature 451:1098–1102. Number: 7182 publisher: Nature Publishing Group.

Sloss, A. N., and S. Gustafson. 2020. 2019 Evolutionary Algorithms Review, pages 307–344. Springer International Publishing, Cham.

Tromp, J. 2016. The number of legal go positions. Pages 183–190 *in* Computers and Games. Springer International Publishing, Cham.

Virtanen, P., R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, d. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, and P. van Mulbregt. 2020. Scipy 1.0: fundamental algorithms for scientific computing in python. Nature Methods 17:261–272. Number: 3 publisher: Nature Publishing Group.

Viswanathan, G., S. Buldyrev, S. Havlin, M. Luz, E. P. Raposo, and H. Stanley. 1999. Optimizing the success of random searches. Nature 401:911–4.

Viswanathan, G., E. Raposo, and M. da Luz. 2008. Lévy flights and superdiffusion in the context of biological encounters and random searches. Physics of Life Reviews 5:133–150.

Walraet, M., and J. Tromp. 2016. A googolplex of go games. Pages 191–201 *in* Computers and Games. Springer International Publishing, Cham.

Werner, E. E., and D. J. Hall. 1974. Optimal foraging and the size selection of prey by the bluegill sunfish (lepomis macrochirus). Ecology 55:1042–1052. Publisher: Ecological Society of America.

Wes McKinney. 2010. Data Structures for Statistical Computing in Python. Pages 56 – 61 *in* Stéfan van der Walt and Jarrod Millman, eds. Proceedings of the 9th Python in Science Conference.

Wiser, M. J., and R. E. Lenski. 2015. A comparison of methods to measure fitness in escherichia coli. PLoS ONE 10:e0126210.

Wosniack, M. E., M. C. Santos, E. P. Raposo, G. M. Viswanathan, and M. G. E. da Luz. 2017. The evolutionary origins of Lévy walk foraging. PLOS Computational Biology 13:1–31.

Zhou, H., S. Yang, and S. Zhang. 2018. Conformable derivative approach to anomalous diffusion. Physica A: Statistical Mechanics and its Applications 491:1001–1013.

Figure legends

Figure 1: Examples of Lévy dust environments. Panels A-C, $E_{997\times997}$ with $n^2 \cdot 10^{-1}$ (rounded down to the nearest integer) Lévy dust (LD) distributed resources with $u = 1$, 2, and 3, respectively. Panels *D-F*, histograms of the distributions used to produce panels *A-C*. We used 100-fold more resources than in our simulations to better illustrate the spatial nature of each exponent, and log-transformed the distributions to elucidate differences in shape; longer move lengths decrease in frequency with increasing u. Note: the y-axes of panels *D-F* are normalized frequencies.

Figure 2: Modelled assortment and syngamous offspring production. Following selection (see text), one of two halves of each randomly rearranged parental genomes (G, a finite, unordered list of move lengths l) is assigned through independent assortment as the gamete to form one of four possible offspring genomes (indicated by the '−−', '· · · ', '—', and '-··' lines). An offspring's genome is thus composed of a random sample of move lengths inherited equally from both selected parental genomes.

Figure 3: Evolution of foraging behaviour in response to LD $u = 2.0$ resources (fig 1B). The labelled red points highlight changes in the move-length distributions over generational time. Panel *A* shows the change in mean fitness (AOL = average number of resources found over a lifespan) starting from a uniform distribution. Panels *AI, AII, AIII* are frequency plots of the ancestral and derived distributions corresponding with the labelled points in *A*, occurring at the 0th, 100th, and final generation. Panel *B* shows the change in mean fitness starting from an all-ones distribution. Panels *BI, BII, BIII* are frequency plots of the ancestral and derived distributions corresponding with the labelled points in *B*, occurring at the 0th, 2000th, and final generation. The y-axis breaks in panels *AIII*, *BII*, and *BIII* highlight the secondary mode of the distributions.

Figure 4: Geometric distribution-based model fitted to the derived distributions. Panel *A* is the move-length probability distribution averaged over all individuals. Panel *B* contains the super-imposed CDF of each individual in comparison to the best-fit CDF obtained from equation (3) with $p_{primary} = 0.9875$ and $p_{geometric} = 0.33$. Panels *C-E* are walks derived from LD environments with $u = 1, 2, 3$, respectively, whereas panel *F* shows a walk from the model of best-fit.

Figure 5: Impact of resource density and distribution on the distribution of realized move lengths pulled from a fixed derived distribution shown in figure 3, panel *AIII*. Realized distributions of move lengths over each tested environment – ballistic ($u = 1.0$), Lévy ($u = 2.0$), and Brownian ($u = 3.0$) distributed resources – and with the tested density of $10^{-3} \cdot n^2$ resources in addition to a 10-fold factor above and below. Each distribution was generated from the truncated moves of a 5 walks, each with lifespan $\lambda = n^2$ steps. Note: move lengths are on a logarithmic scale.

Figure 6: Number of generations required for the derived populations to displace the Lévy populations from the competition simulations. Panel *A* is lifespans $\lambda = n^2$ and panel *B* is lifespans $\lambda = 20 \cdot n^2$.

Video S1: Animation of figure 3A, evolution of distributions starting with random uniform move lengths, where successive frames are the successive generations saved within our digital fossil record. The histogram represents the concatenated distributions of all individuals.

Video S2: Animation of figure 3B, evolution of distributions starting with 'all-one' move lengths, where successive frames are the successive generations saved within our digital fossil record. The histogram represents the concatenated distributions of all individuals.

Video S3: Animations of individual walks over 10^4 steps. Panel *A* is a walk from one of the derived distributions, panel *B* is the model fitted walk (same as fig 4F), and panel *C* is a Lévy Walk.

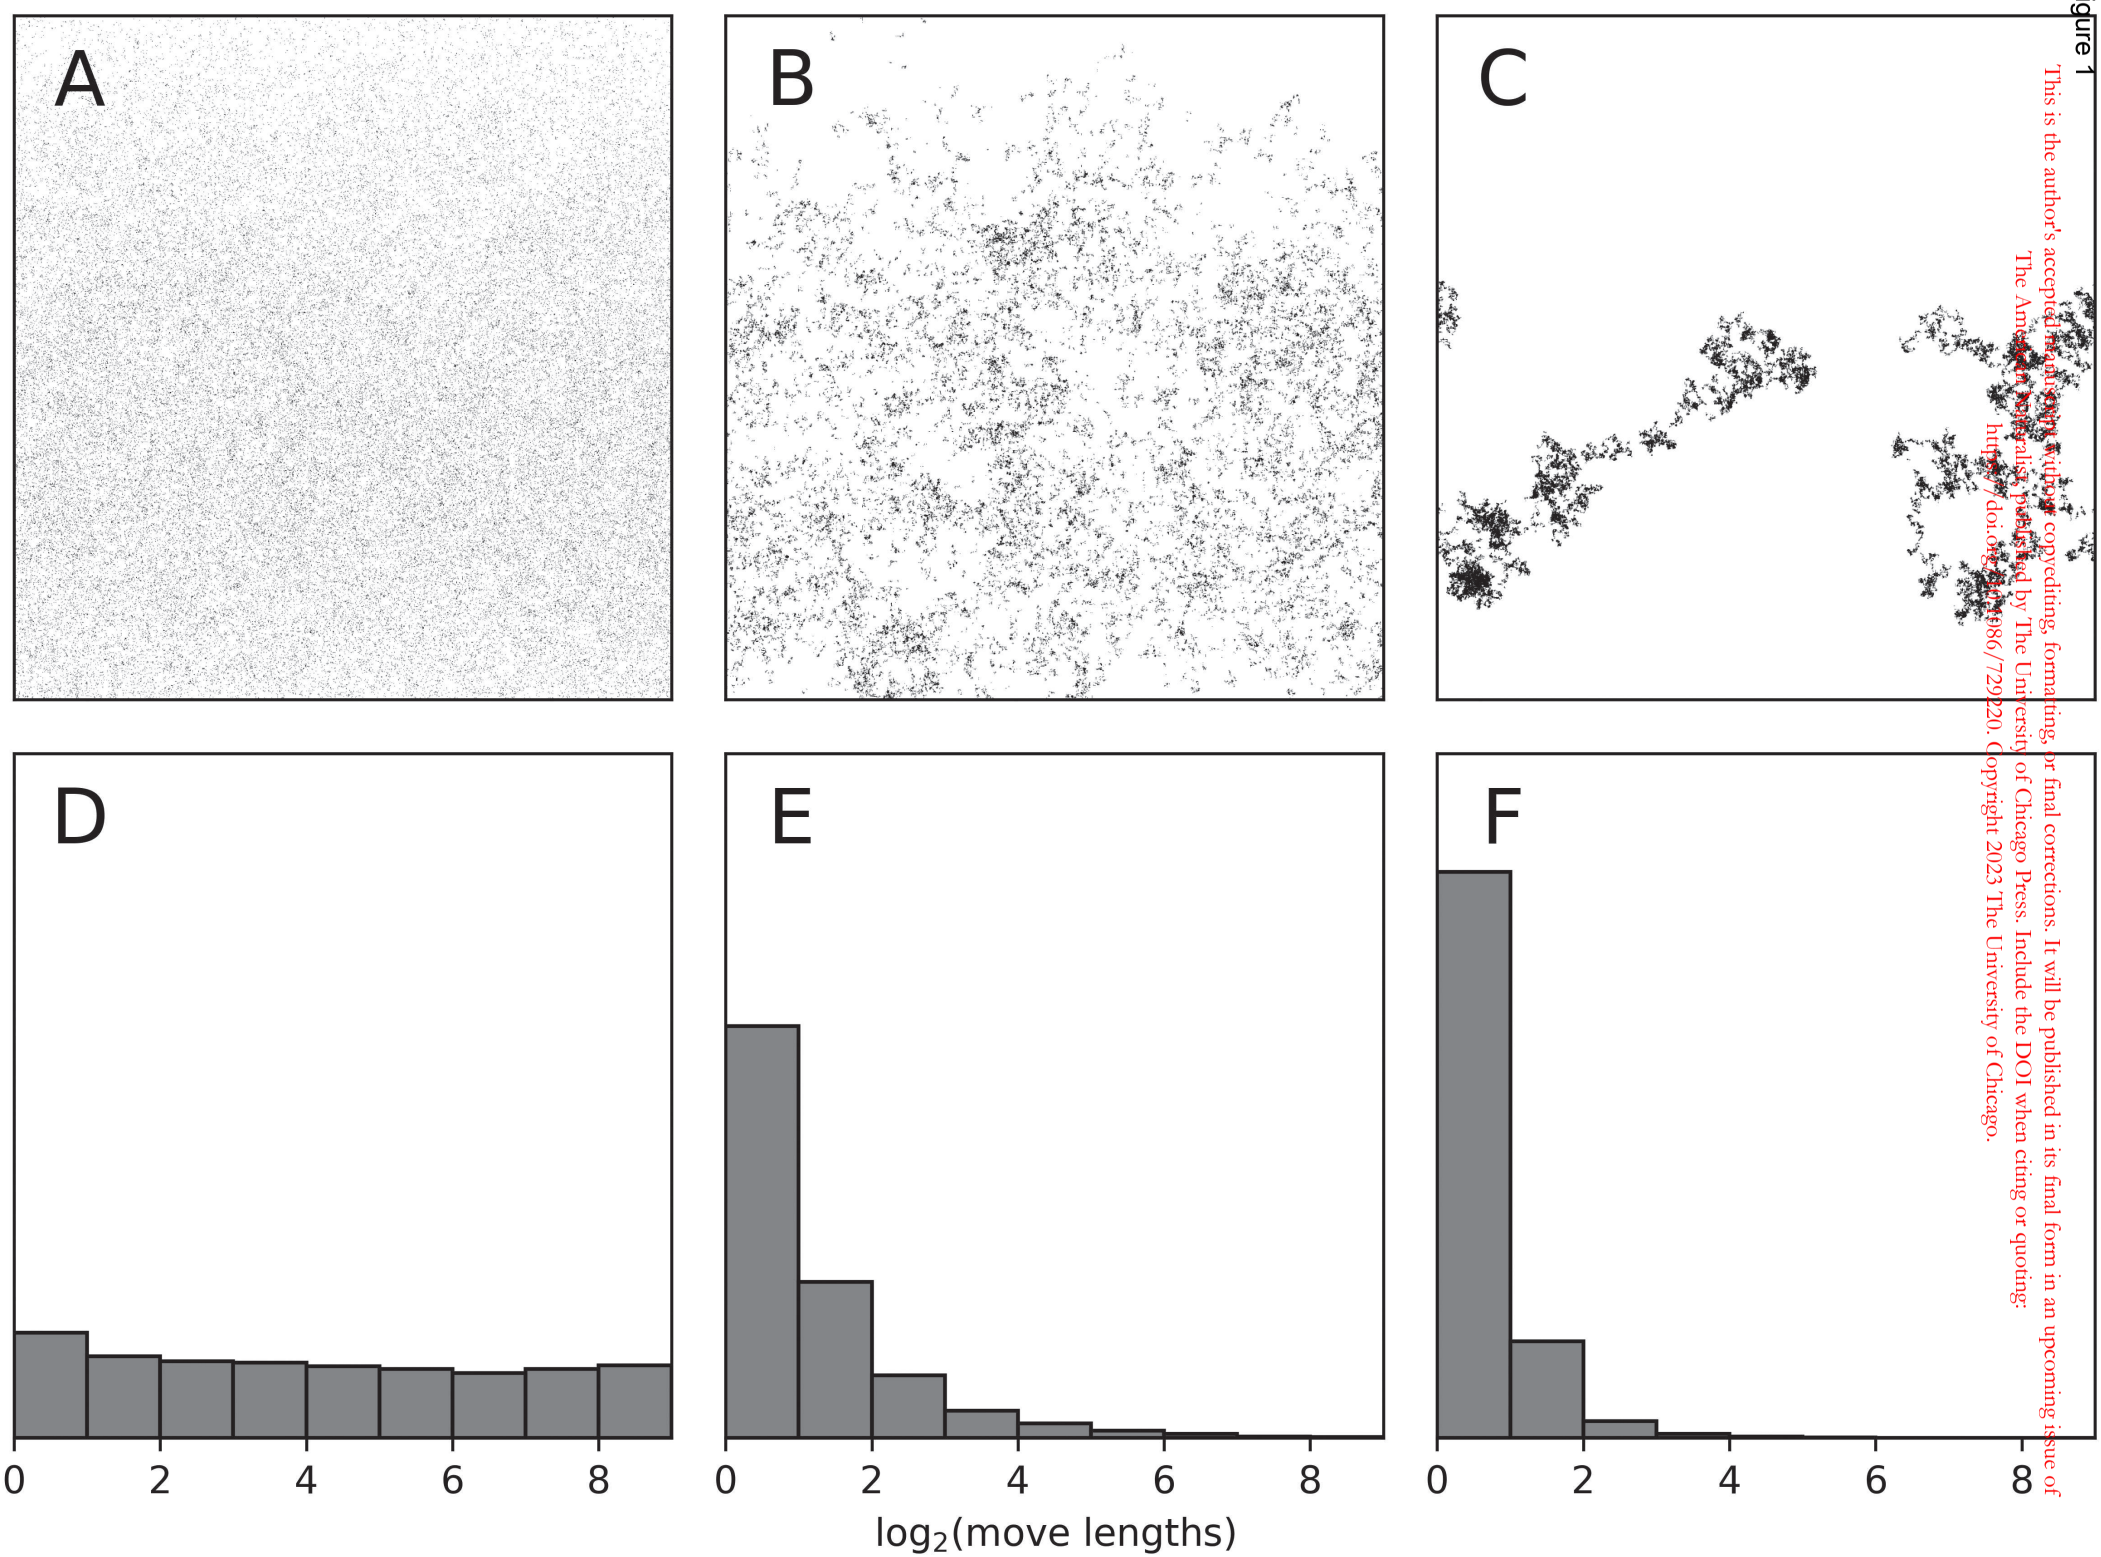

Figure 1

Figure 2

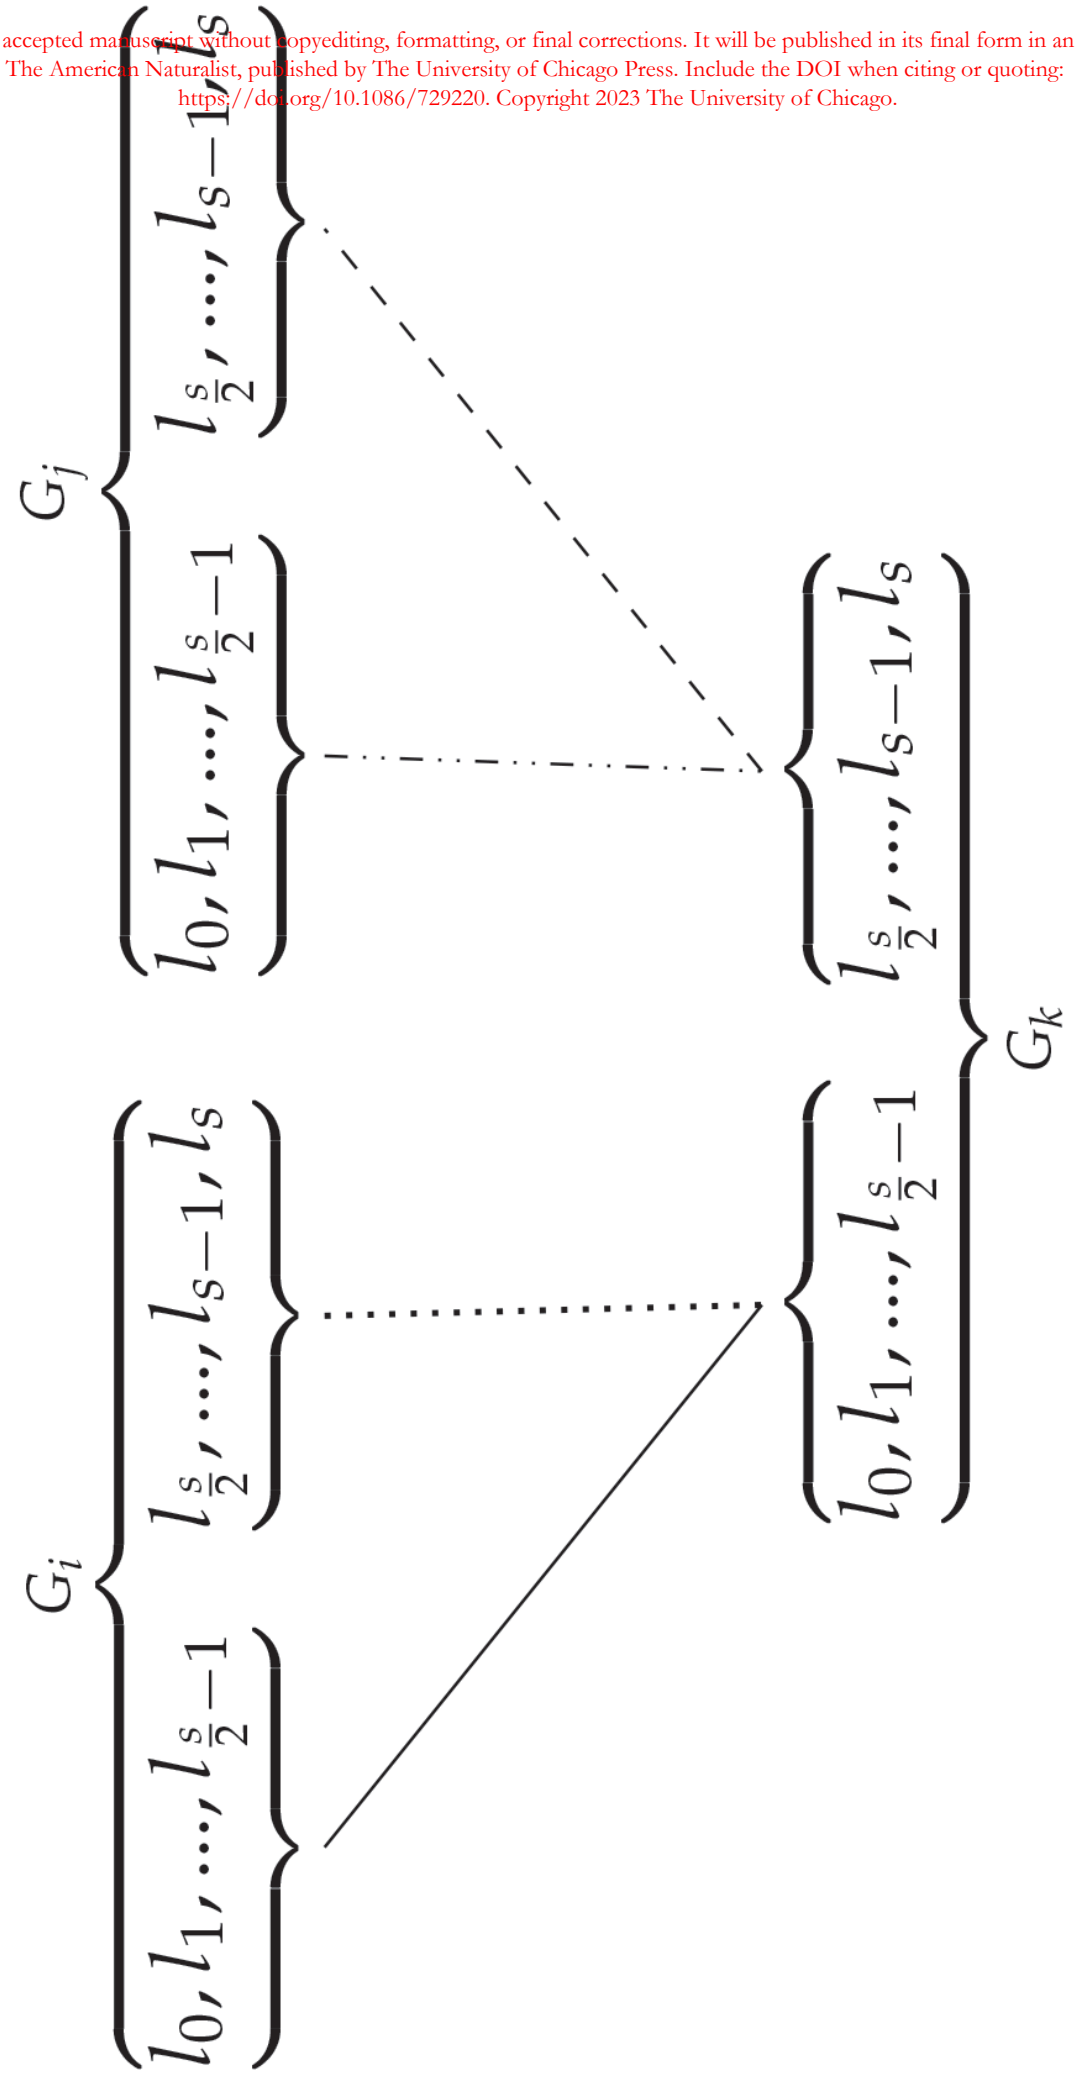

Figure 3

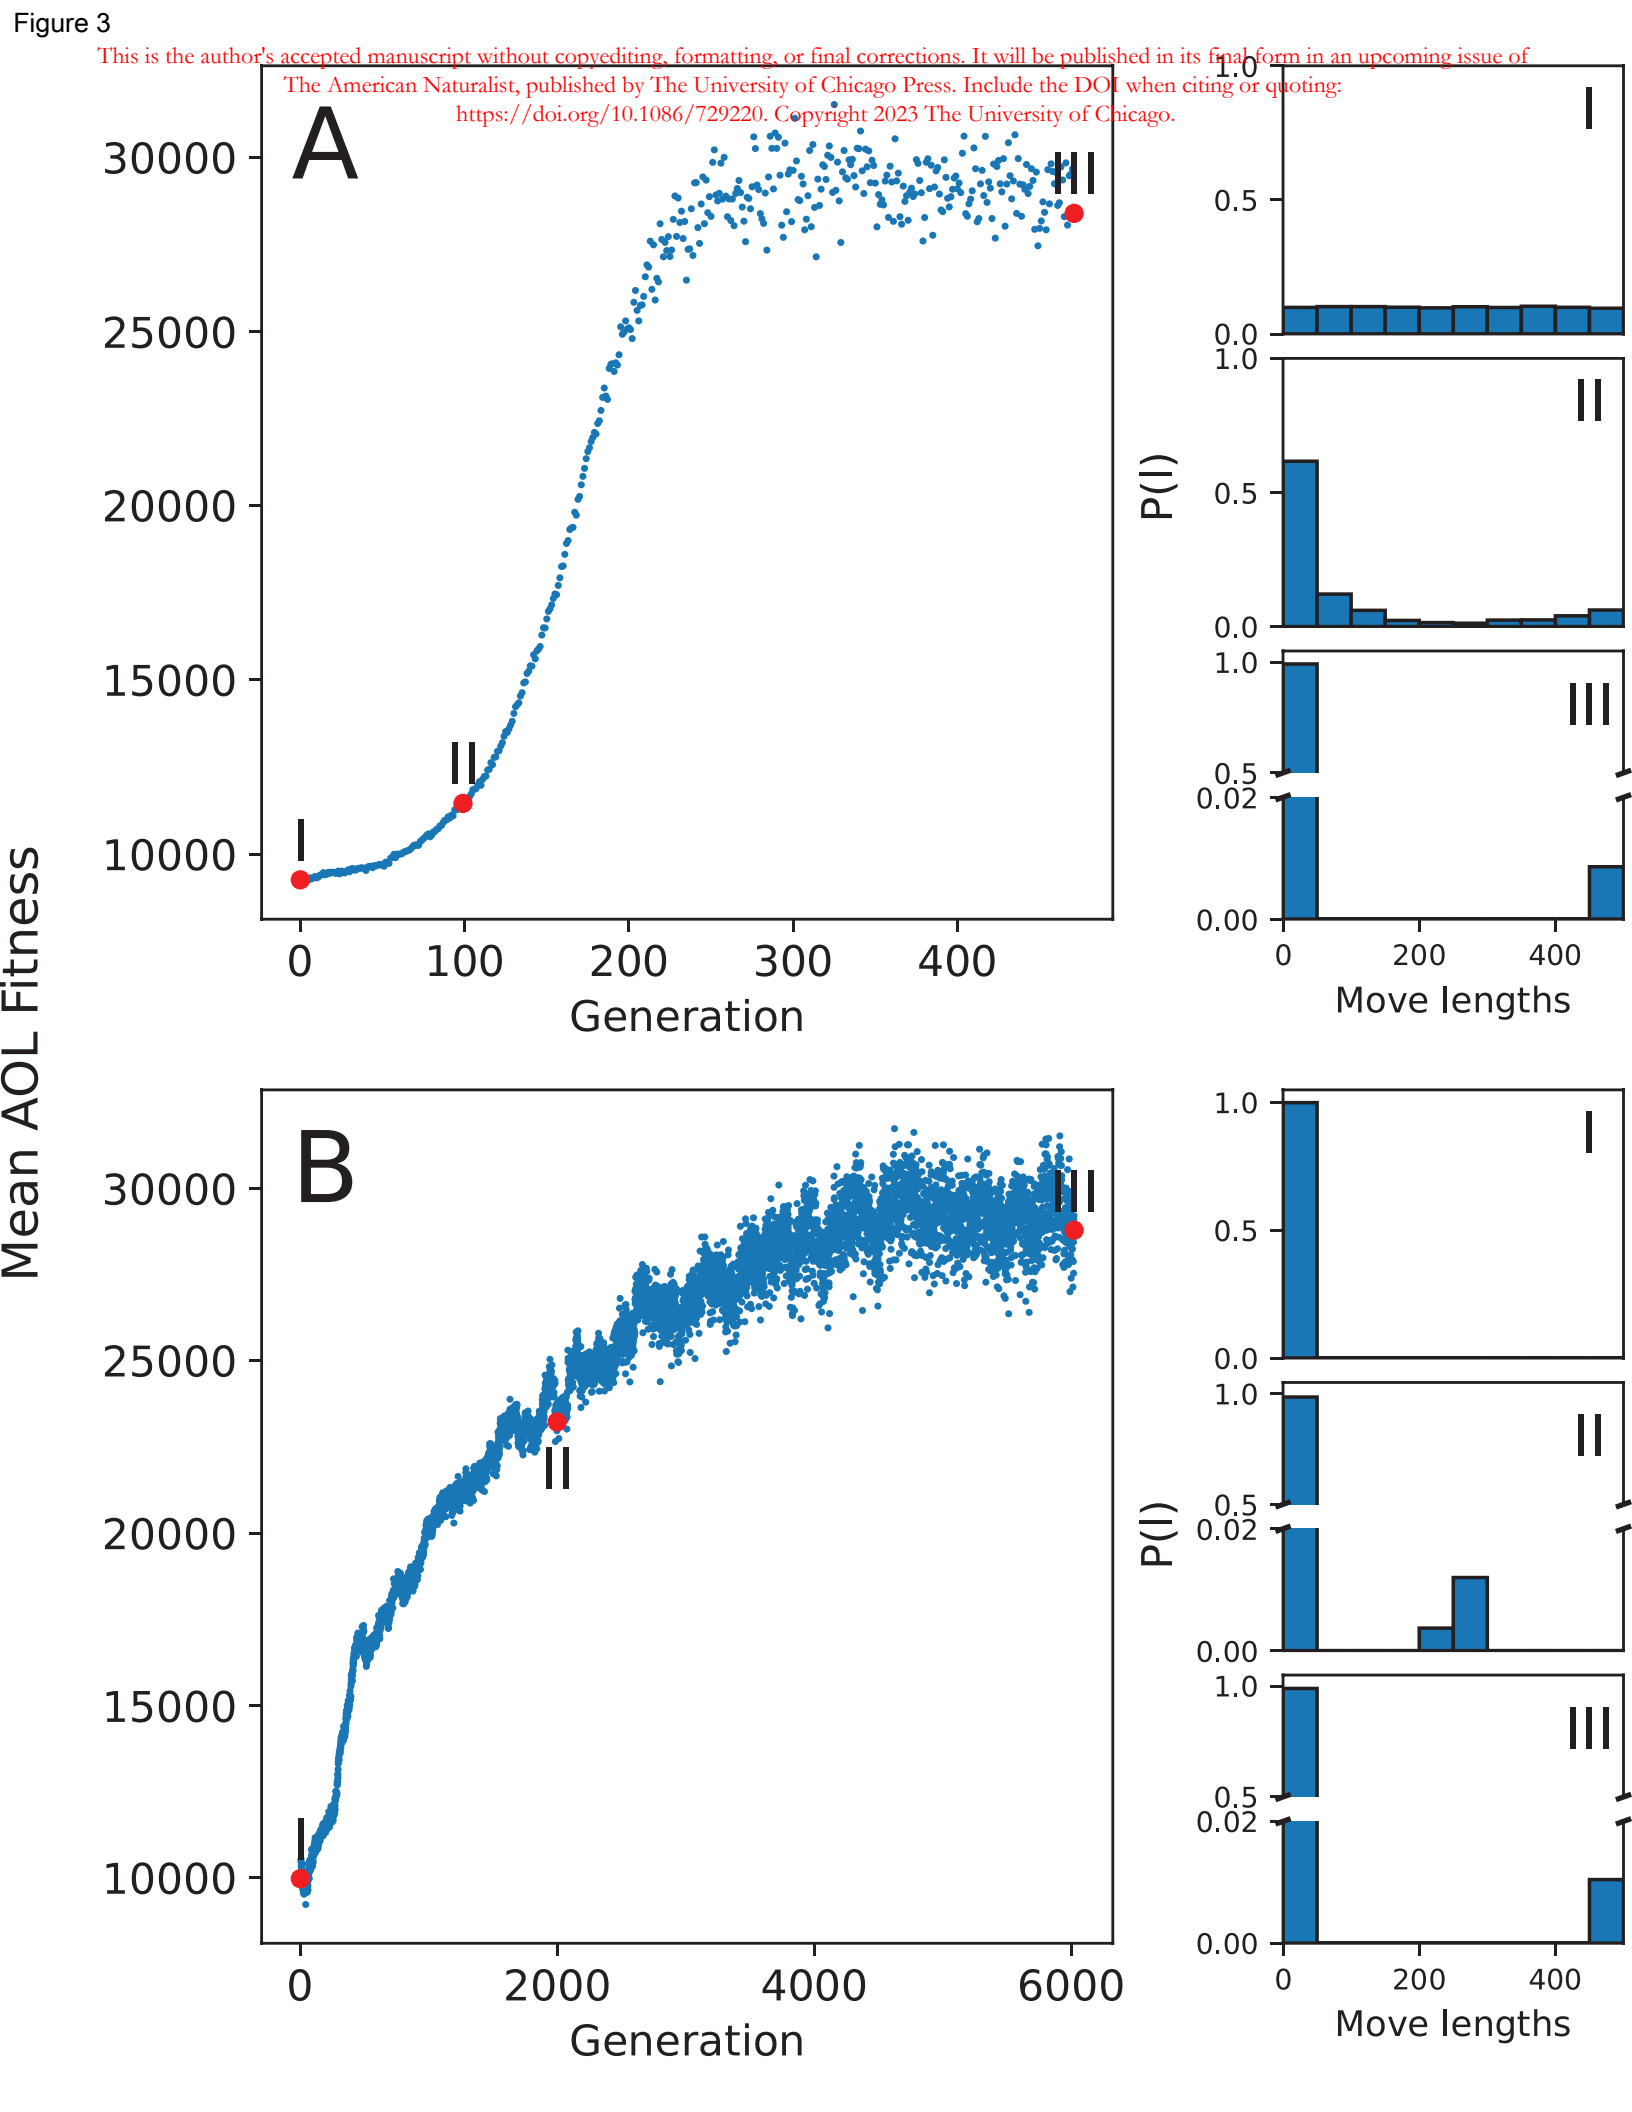

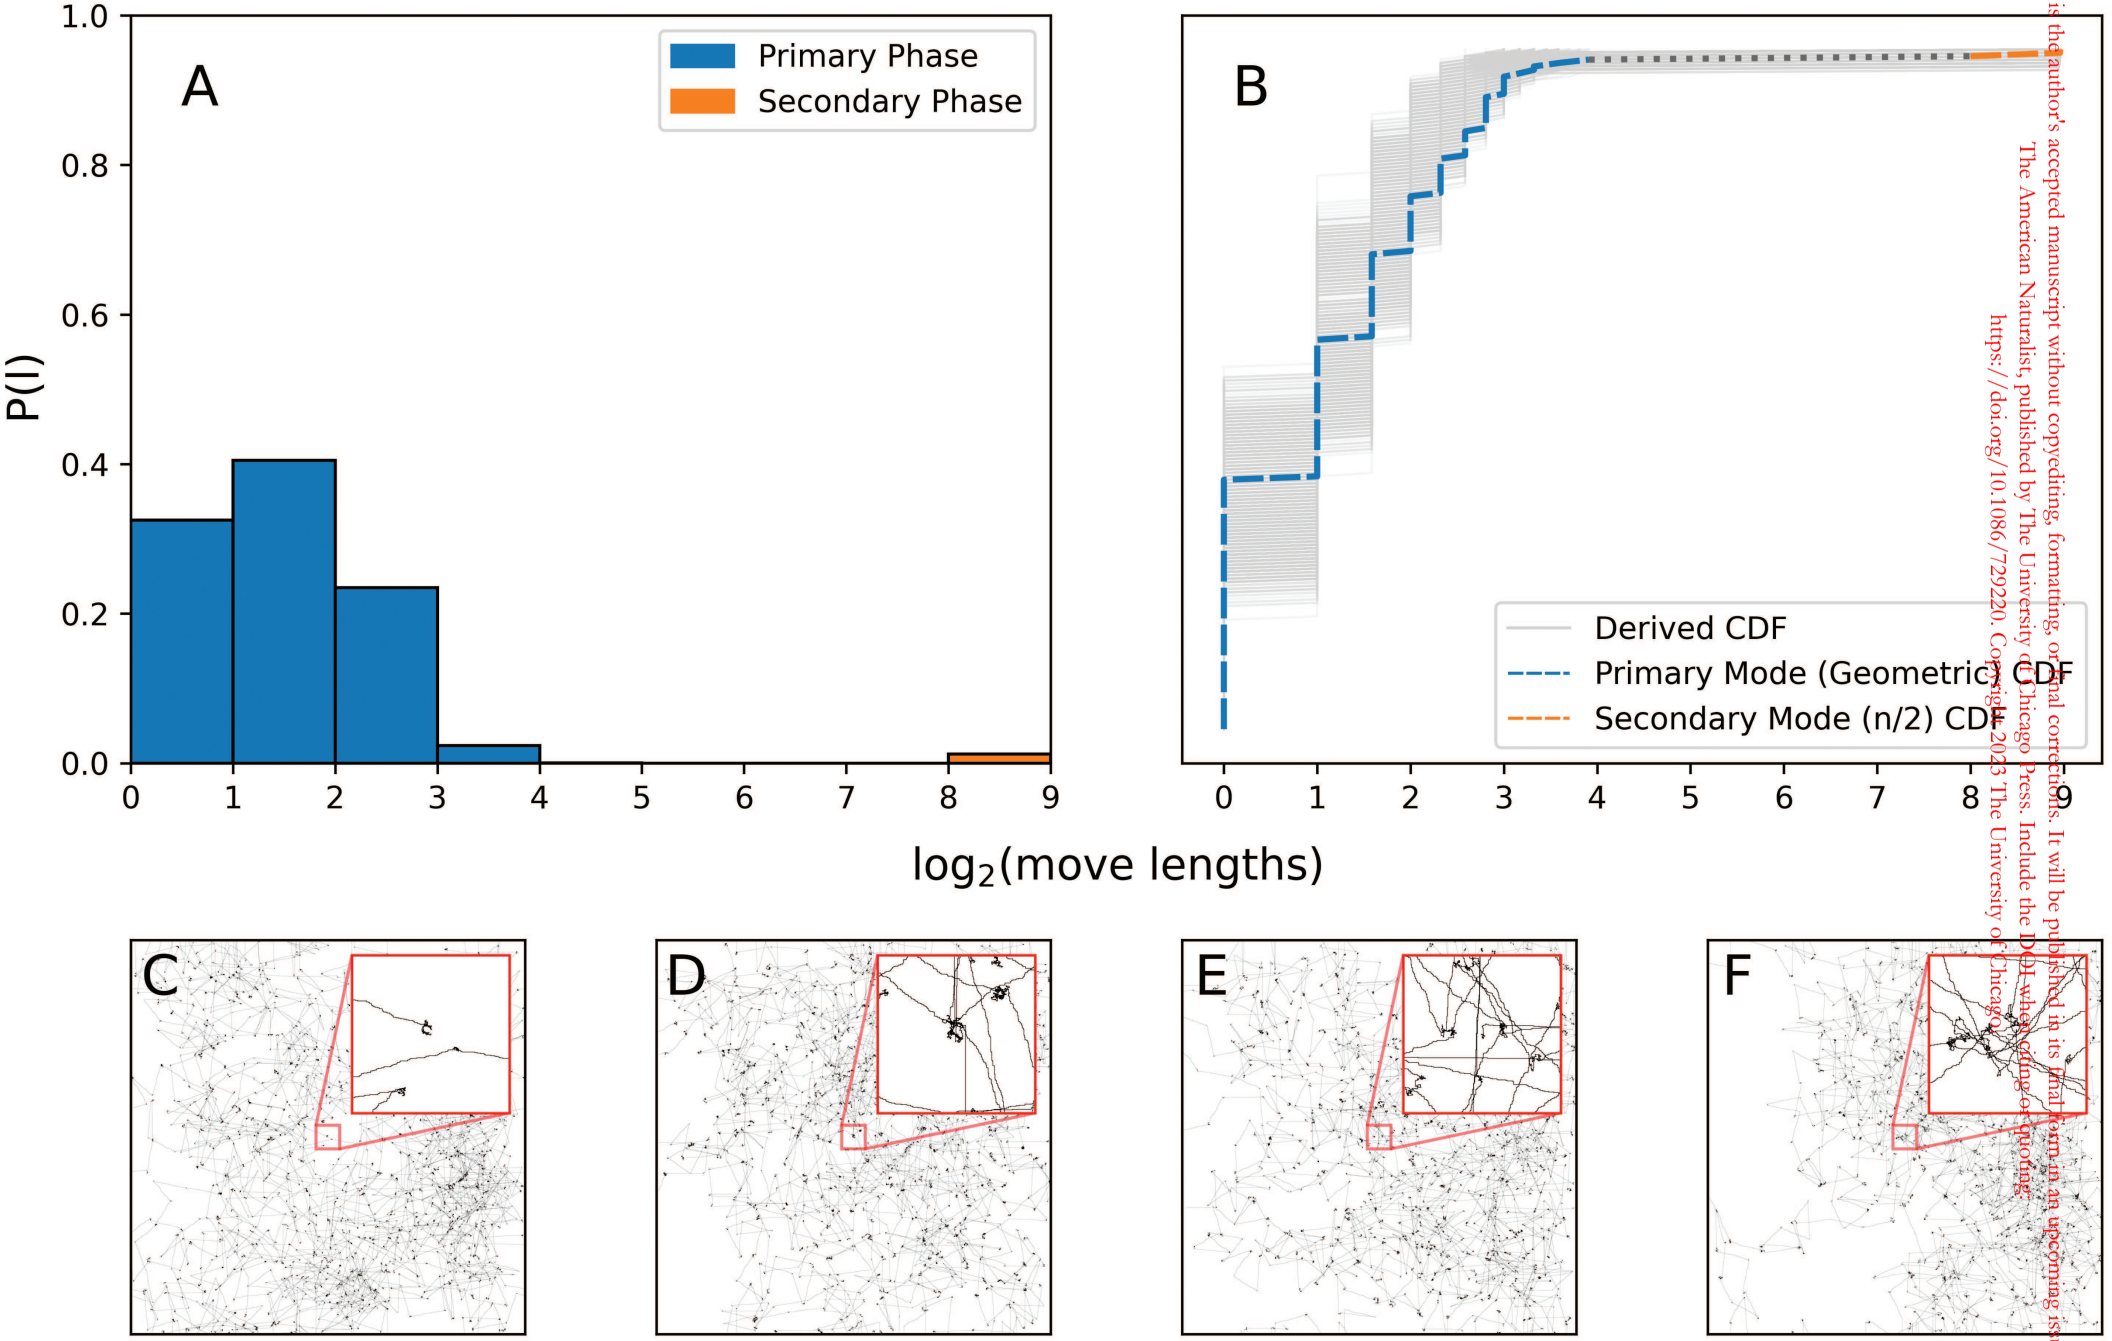

Figure 4

Figure 5

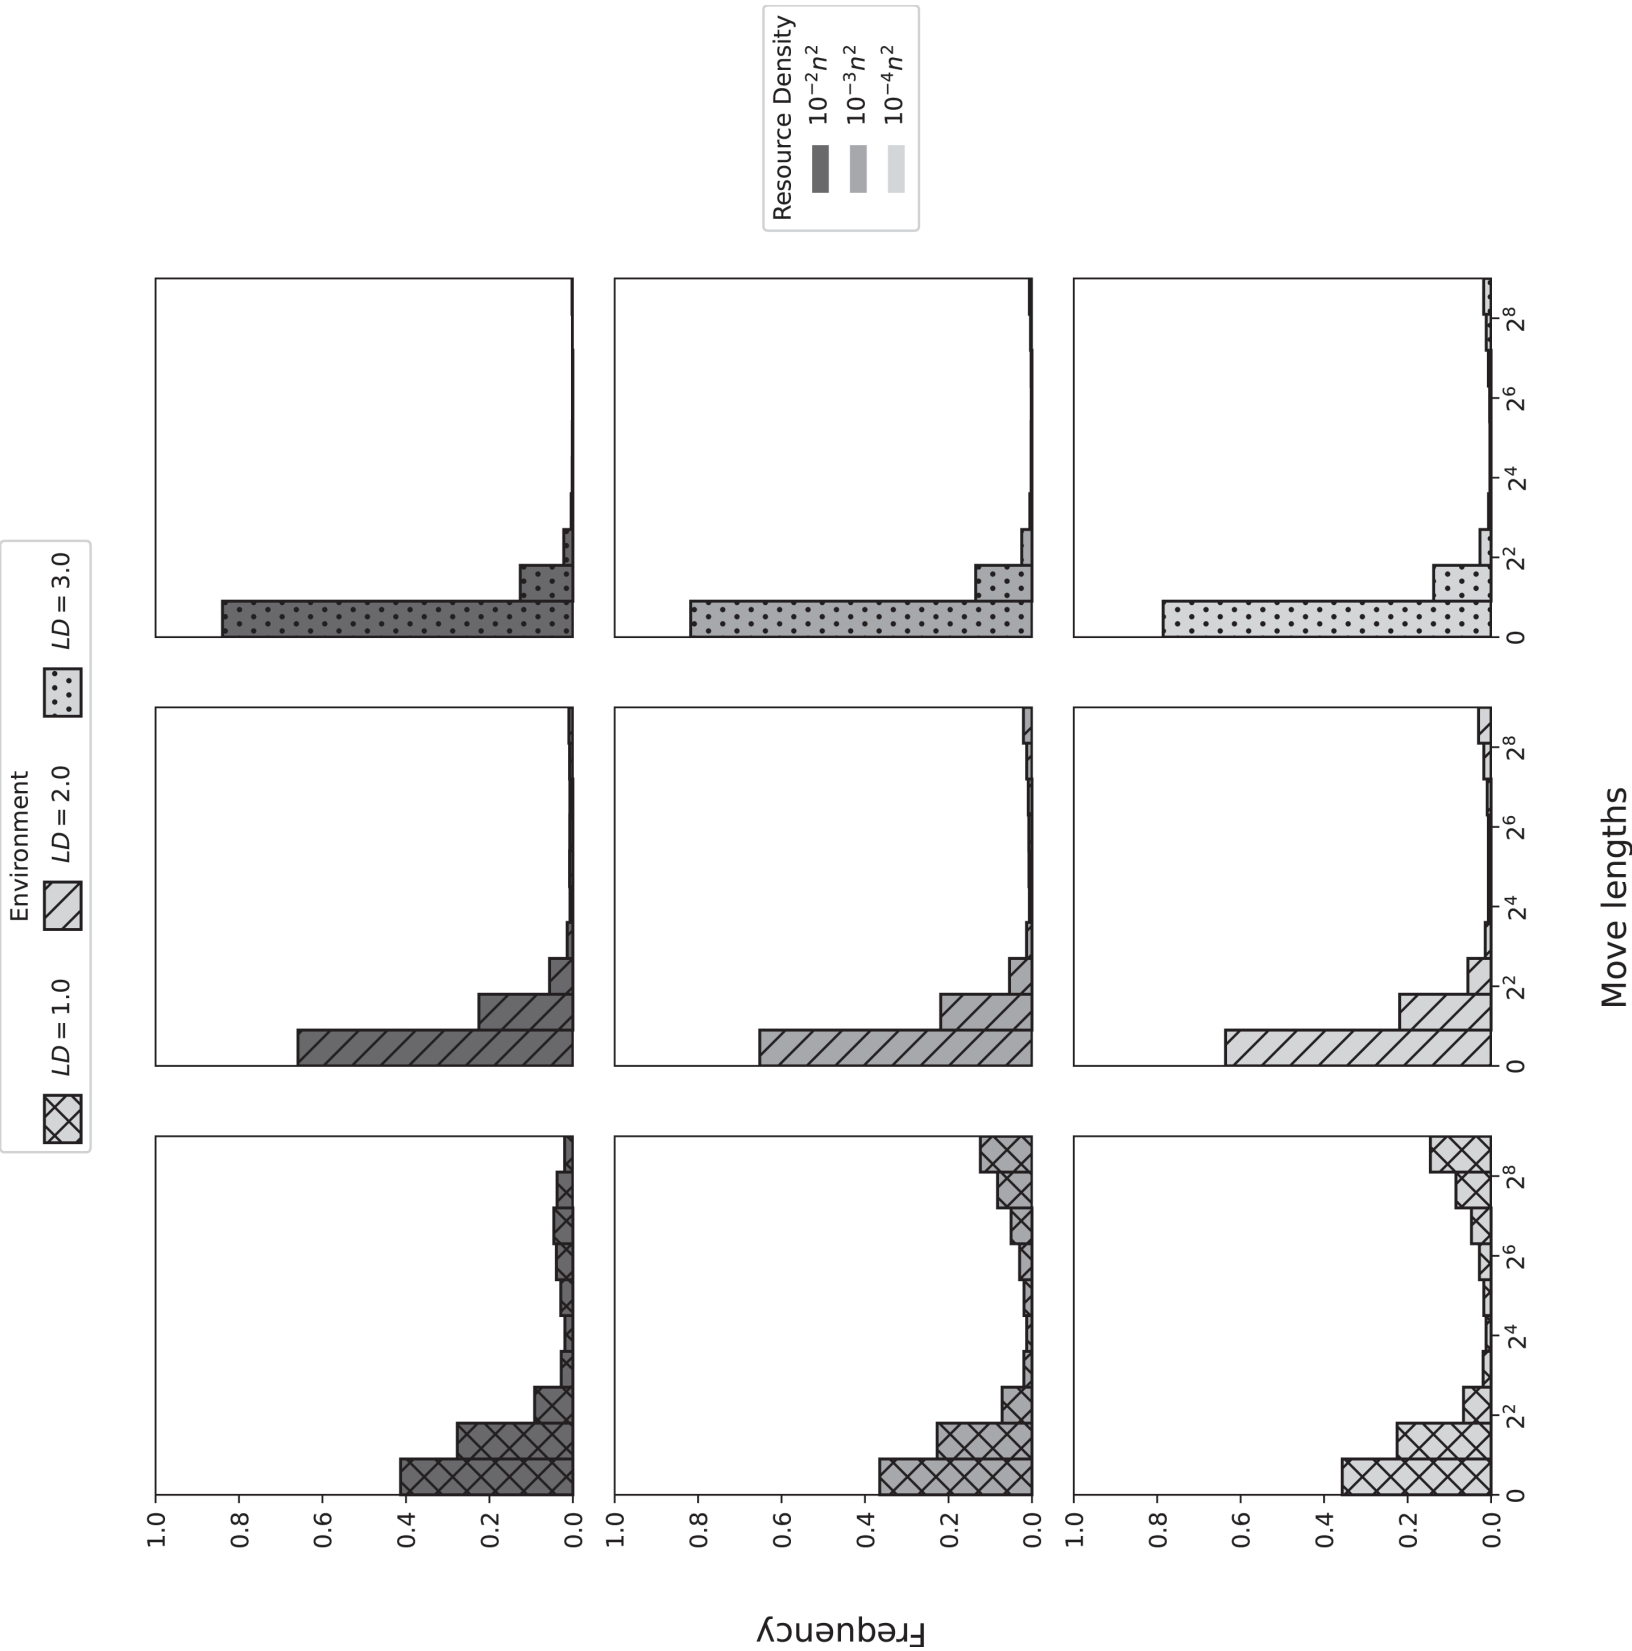

Figure 6

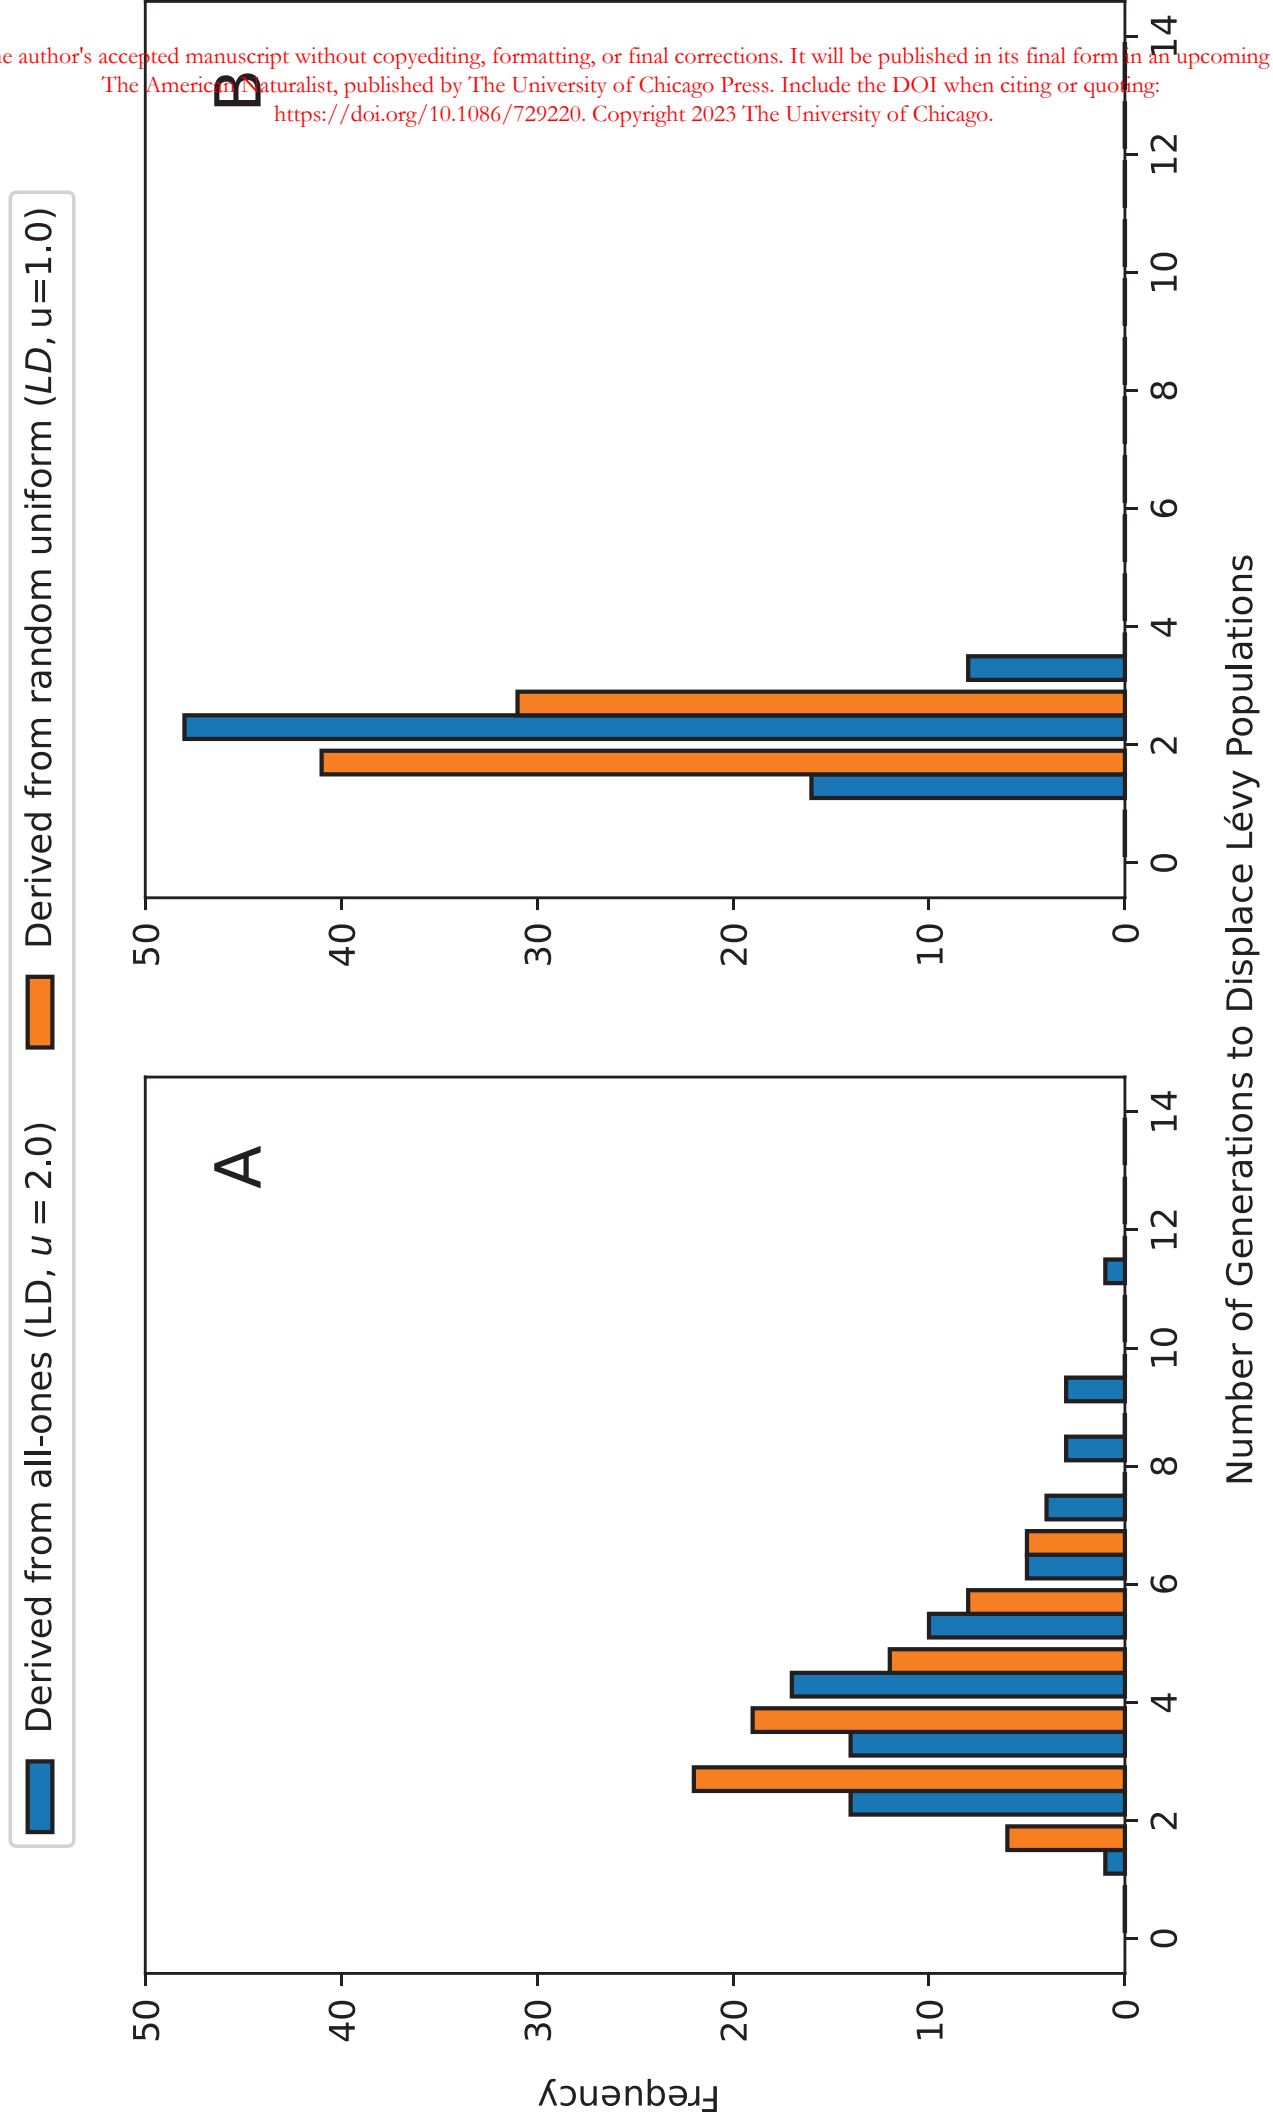

This is the author's accepted manuscript without copyediting, formatting, or final corrections. It will be published in its final form in an upcoming issue of The American Naturalist, published by The University of Chicago Press. Include the DOI when citing or quoting: https://doi.org/10.1086/729220.

Online Supplement:

Intermittent search, not strict Lévy flight, evolves under relaxed foraging distribution constraints

The American Naturalist

Winston Campeau[1,*]

Andrew M. Simons[1]

Brett Stevens[2]

1. Department of Biology, Carleton University, Ottawa, Ontario K1S 5B6, Canada.

2. School of Mathematics and Statistics, Carleton University, Ottawa, Ontario K1S 5B6, Canada.

∗ Corresponding author; e-mail: winston.campeau@carleton.ca.

Supplementary Figures & Results

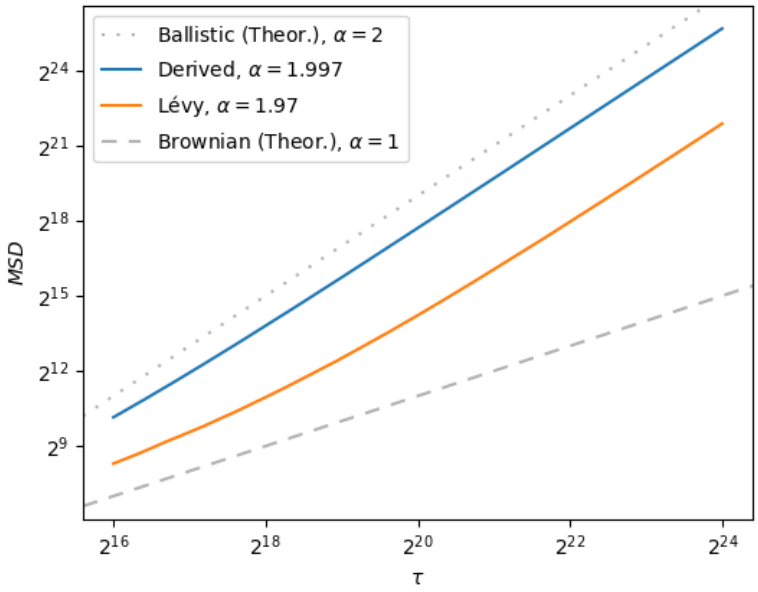

Figure S1: Mean squared displacement (MSD), a measure of diffusivity, as a function of lag time τ, averaged over 10^4 walks, each with a lifespan of $\lambda = 2^{24}$, for Lévy and derived distributions. The MSD was sampled every 2^8 steps, but the plot begins $\tau = 2^{16}$, where the sampling rate is higher between log-intervals, for better visual comparisons of the trends. Note: the axes are on logarithmic scales.

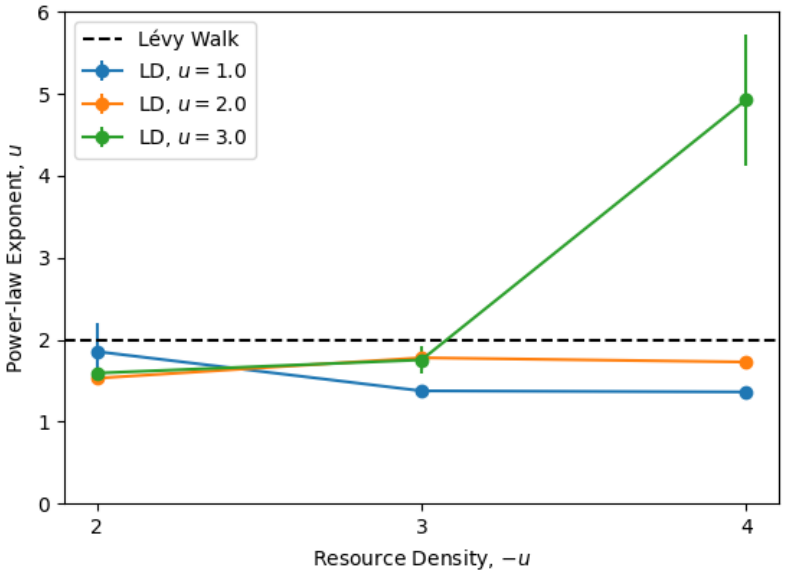

Figure S2: Power-law analysis of the realized distributions of move lengths over different resource distributions and densities. Each point is the average of five trials from an individual with lifespan $\lambda = n^2$. Error bars are the standard error of the mean power-law exponent, and are otherwise not visible if the error was small.

Competition Simulations

We computed an additional set of competition simulations with very small lifespans $\lambda = 0.1 \cdot n^2$, using the derived distributions from the random uniform starting condition, only. Here, we also provided the Lévy walks with access to strictly a G of size 2000. The derived distributions won 35 of the 36 simulations. The median number of generations needed to displace the Lévy walks was 8, with a minimum of 1, and maximum of 26. In the one instance were the Lévy walks succeeded, the population size was 10 and the simulation elapsed over 11 generations, as compared to the other two replicates where the derived distributions displaced the Lévy walks in only four generations. Due to the small lifespan and population size, we ascribe this singular loss as a

result of random chance.

Whereas the average probability of drawing a move length of $n/2$ between the two modes was approximately 1% among all derived populations, variation within each population differed, although was similar between the environments. The derived distributions evolved in ballistic environments (LD, $u = 1.0$) had one or zero individuals with no secondary mode, eight or nine individuals with no secondary mode in Lévy environments (LD, $u = 2.0$), and four individuals with no secondary modes in Brownian environments (LD, $u = 3.0$). Our competition simulations included the populations with the second largest difference in missing secondary modes; zero from the random uniform distributions evolved under ballistic environments, and eight from the all-ones distributions evolved under Lévy environments. To determine whether the proportion of strict primary modes was adaptive, we first compared the generational time required to displace the Lévy populations among the two derived populations in their original and opposing environments (fig. S3). In either case, the derived population with only bimodal walkers required fewer generations to displace the Lévy populations. We then competed the derived population with eight primary-mode walkers against itself in each environment type, with all 200 individuals, lifespans of $\lambda = 20 \cdot n^2$, and patch density $n^2 \cdot 10^{-3}$, to determine whether the proportion of strict primary modes were adaptive. The bimodal individuals completely displaced the strict primary-mode individuals in each instance, with an approximate overall increase of 7% in fitness after a single generation, 9% after four generations, and 10% after eight generations.

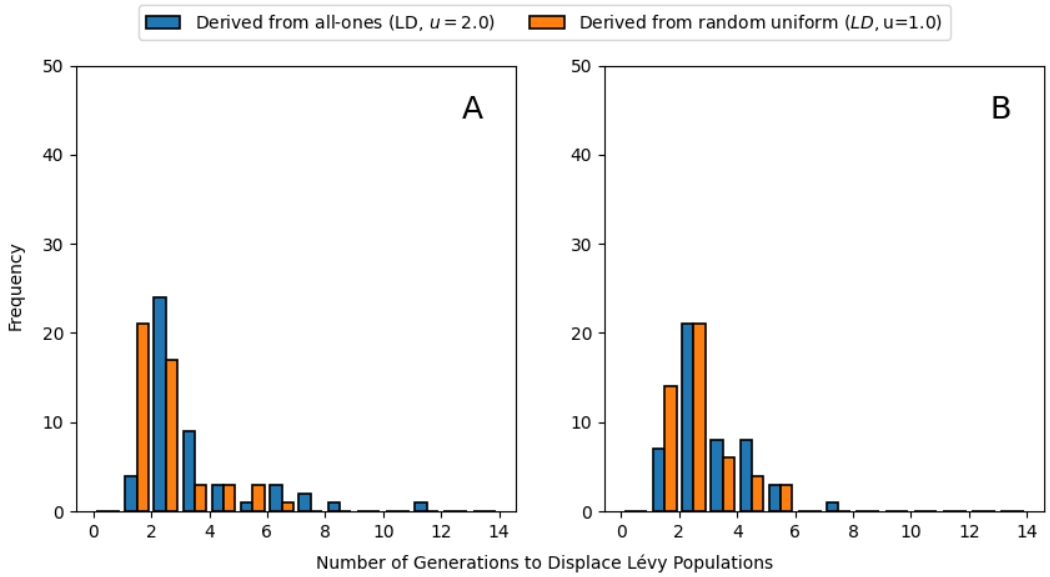

Figure S3: The number of generations required for the derived populations to displace the Lévy populations in their original and opposing environments. Panel A is ballistic environments (LD, $u = 1.0$) and panel B is Lévy environments (LD, $u = 2.0$).

Additional Methods

Here, we provide detailed descriptions for all of the methodologies we utilized in our simulations. We also discuss the "harmonics" from the testing stages of our simulations, and cover reasoning behind the utility of each method. We noticed that selection would often result in individuals with move lengths which were divisors of n, and that the frequency of a divisor increased as it approached $n/2$, where $n/2$ was always the most frequent divisor (besides $n = 1$). We wanted to ascertain the behaviour of these harmonics, and developed several methods to probe their behaviour. The following section summarizes the choices which were most effective at explaining the harmonics, and which were most consistent with previous methodology. But as we will see, the frequency of $n/2$ move lengths is a feature of the dominant evolved distribution, so its appearance is due to more than these causes.

Each method outlined here resulted in reduced selection for harmonics, but in combination were still not sufficient to completely explain them. In fact, the final method which most effectively accounted for persistent harmonics, was our model's method of recombination. We had previously only used our proxy for independent assortment to generate offspring genomes., without resampling them beforehand. This allowed specific move lengths to persist in the "alleles" and would often result in pre-mature convergence and/or punctuated equilibria. Including a resampling of each genome before assortment resolved the phenomenon of persistent harmonics and pre-mature convergence (at least, up to the derived results presented in this paper). This result is visualized nicely in video S2 where a divisor of n transitions from $l = 1$ to $l = n/2$.

Prime and Larger Dimensions

Say there is an environment $E_{n\times n}$, with a singular re-visitable resource located at $e_{i,j}$. If there is an individual which is restricted to a single movelength, which is also a divisor of n, then that individual is restricted to grid-movement. The grid has $(n/l)^2$ quadrants of size $l \times l$, where each edge of a quadrant is shared with its four neighbours. If the resource located at $e_{i,j}$ happens to be on the edge of any quadrant, then the individual will return to that same resource with very high probability. Conversely, if $e_{i,j}$ does not fall on the edge of a quadrant, then the individual is certain to never encounter the resource (barring the base case of $l = 1$). If a grid intersects a resource, the search efficiency is unusually high because of the restricted number of visitable locations. For example, when $l = n/2$, each edge is shared twice between neighbouring quadrants, each vertex contains the intersection of every quadrant, and there are $4(n-1)$ visitable locations on the grid. The probability of selection resulting in grid-movement is contingent on the population size, number of resources, and the size and divisors of n. A greater population size, or number of resources, increases the likelihood that one of the individuals' grids will intersect with a resource. Environments with larger n decrease the efficacy of a grid pattern by increasing the distance between resources (especially if resources are clumped together). Lastly, if n is prime, the efficacy of grid-movement is altered as it will continually shift; thus, losing the property of

extreme success or failure due to chance intersections with resources.

Whereas we acknowledge the potential efficacy of grid-movement, we believe it is a highly circumstantial result. So, each individual is assigned a unique environment, with larger and prime n (specifically, $n = 997$ for the final results), to reduce the probability of a grid intersecting resources by chance.

Cycling Resource Positions

If grid-movement is successful simply due to chance intersections with resources, then an additional method to control for those chance intersections is cycling resource positions. We create a queue of resources (Fig 2). We generate $2n^2 \cdot 10^{-3}$ (rounded down to the nearest integer) Lévy dust resources, where the first half are distributed onto the environment, and the second half are stored in a queue. Once a resource is located by a individual, it is removed from the environment, and its position is stored at the beginning of the queue. The position at end of the queue is then added to the environment. Queued resources maintain the property of re-visitability, the governing powerlaw exponent, and reduces the efficacy of grid-movement by cycling resource locations. Below is a simple illustration of how resources can be cycled using a queue of resource locations:

We did pursue this methodology as it was not as consistent with previous methodology, whereas procedural-generation was consistent and provided the same solution of increased environmental variance.

Procedurally-generated Environments

Many simulations utilize periodic boundary conditions to approximate larger environments, including research on Lévy flight foraging. However, revisiting the same environment invites the possibility of selection for features unique to a particular environment. Larger environments

$$Initial\ Resources = [e_{0,0}, e^{*}_{1,1}, e_{2,2}]$$

$$Queued\ Resources = [e_{0,2}, e_{1,2}, e^{@}_{2,0}]$$

$$\underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1^{*} & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{Initial\ Environment} \rightarrow \underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 1^{@} & 0 & 1 \end{pmatrix}}_{Updated\ Environment}$$

$$Updated\ Queue = [e^{*}_{1,1}, e_{0,2}, e_{1,2}]$$

Figure S4: An individual encounters a resource located at $e_{1,1}$ (indicated by *) in the initial environment. The resource is removed from the environment and placed at the beginning of the queue. The environment is then updated by adding a resource to the location at the end of the queue $e_{2,0}$ (indicated by @), and the location is subsequently removed from the queue.

can also be achieved with procedural generation, removing the need for periodic boundary conditions on individual movement. Before the main simulation begins, a library of unique environments is generated (2^{10} by default), and each individual is assigned a random starting environment from the library. Each environment has the same distribution and density of resources, as described in the first subsection of our methods. New environments are appended when an individual crosses over the edge of its current environment, and the relative locations are saved in an expanding dictionary. Our method of procedural generation increases environmental variability and reduces the possibility for selection of artifacts due to revisiting the same environment.

In our simulations we use lifespans which theoretically allow each individual to span a single environment 20 times. Thus, traversing a single environment 20 times, which also has sparsely

distributed resources, has a stronger possibility for selection resulting in adaptations for features which occur by chance. Note that although we generate a library of 2^{10} environments, this is likely "overkill", and a realistic minimum library size is the spanning number, i.e., as many environments for as many times an individual can theoretically span a single environment (20 in the case of this paper, or 2^5 if you insist on using powers of two). Also note, a library containing only one environment, is equivalent to a typical model with periodic boundary conditions.

Turning Algorithm

As a grid is a network of rectangles, selection for grid-like movement may be exacerbated by turning angles which are always perpendicular to the axes. Adding to the previous methodology, we employed an alternative and more complex method of movement which more closely resembles selecting a random direction from $[0, 2\pi)$. Upon selecting a move length l, a random selection from the border of all locations exactly l distance away becomes the endpoint (i.e., the ℓ_1 radius). That move length is the sum of displacements, say x, y, such that $|x| + |y| = l$. Whereas resources are simply placed at the end of successive flights, an individual must walk to the endpoint, and this can be achieved by any permutation of the increments in x and y. Instead of choosing arbitrary paths or permutations, we allow for all possible paths. A random selection without replacement is made from a list of vertical and horizontal increments leading towards the chosen endpoint, and this process continues until the list is empty. The probability of visiting any location $e_{i,j}$ with a move length of l, which is the sum of its x and y components, is:

$$\binom{i+j}{i}\binom{x+y-i-j}{x-i}\binom{x+y}{x}^{-1} \tag{S1}$$

i.e., the product of all possible paths from the starting location to $e_{i,j}$, from $e_{i,j}$ to $e_{x,y}$, and the inverse from the starting location to $e_{x,y}$. Note that whereas this method of movement eliminates the bias of choosing an arbitrary path, and concomitantly increases sampling of the environment, the diagonal remains the most probable path (fig. S5).

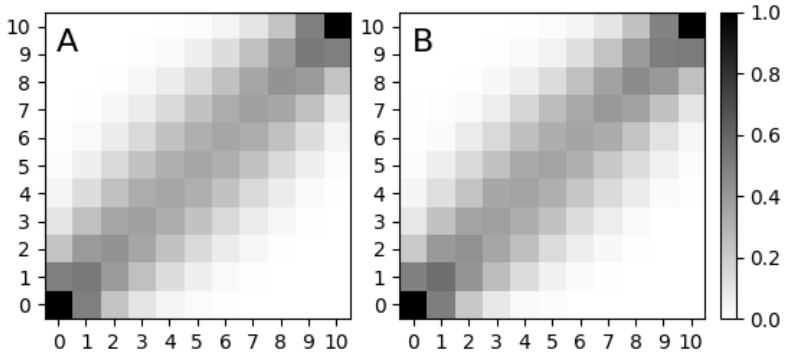

Figure S5: An example of the resultant probability space from our turning algorithm by a move length $l = 20$, from location $[0, 0]$ to $[10, 10]$. Panel *A* is the theoretical probability for visiting each location, as calculated in equation S1. Panel *B* is the probability of visiting each location averaged over 10^3 simulated walks. The scale maps the gradient to the corresponding probabilities.