Ants on a Stick

Collision Invariance, Order Statistics, and Monte Carlo Estimation

Arian Nazarian

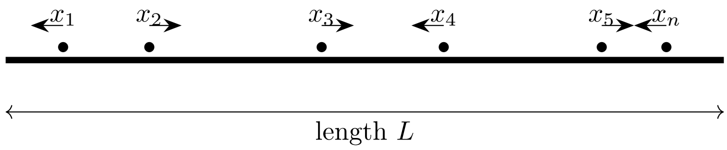

Figure 1: Identical point agents on an interval. Each moves at speed v in one of two directions, reverses direction at collisions, and exits at an endpoint.

Abstract

The ants-on-a-stick puzzle looks like a collision-tracking problem, but equal-speed reversals admit an exact relabeling: for any observable that ignores individual identities, colliding ants may be replaced by noninteracting trajectories that pass through one another. The full exit-time multiset is therefore determined by the directed distances from the initial positions to the endpoints they face.

This reduction turns the puzzle into a compact study of deterministic bounds, algorithms, random inversions, order statistics, stochastic exit processes, extreme values, and Monte Carlo estimation. Completion time is computable in optimal $\Theta(n)$ time even though the physical system may contain $\Theta(n^2)$ collisions. Under fair independent directions, the collision count has exact mean $n(n-1)/8$ and variance $n(n-1)(2n+5)/96$.

On the probability side, a reflection principle identifies the minimal symmetry behind direction-law invariance: if the iid initial-position law is symmetric about the midpoint, then any independent joint law of directions leaves the directed-distance vector iid with the original position law. Uniform positions therefore reduce the entire exit process to uniform order statistics. The normalized inter-exit spacings are Dirichlet, the repeated-trial average completion time converges almost surely to $\frac{n}{n+1}\frac{L}{v}$, the ants-remaining counter is an exact time-inhomogeneous pure-death process, and the terminal order statistics have a Gamma hierarchy whose first case is the reverse-Weibull completion-time limit. A general completion-time formula is also obtained for arbitrary continuous position laws and arbitrary independent direction vectors.

Because the stochastic target law is known exactly, the model gives an auditable Monte Carlo benchmark. Antithetic reflection has correlation $-1/n$ with the original completion-time functional and yields closed-form variance ratios under two different cost conventions. The companion browser piece uses the same deterministic bound, exit process, and randomized motion as distinct visual ingredients.

1 Problem and Model

A standard ants-on-a-stick puzzle places identical ants on a finite stick, gives each ant the same speed, and has it reverse direction whenever it collides with another ant. Once an ant reaches an endpoint, it falls off. The familiar observation is that a head-on collision between two identical equal-speed ants has the same unlabeled effect as allowing the two trajectories to pass through each other and exchanging their labels.[6]

Let the stick be the interval $[0, L]$, with $L > 0$. There are $n \geq 1$ point ants at distinct initial positions

$$0 \leq x_1, \ldots, x_n \leq L,$$

and each has direction

$$d_i \in \{-1, +1\},$$

where -1 means left and $+1$ means right. Every ant moves at the same speed $v > 0$. Collisions are instantaneous and reverse both directions. An ant initialized at an endpoint and directed outward exits at time zero; one initialized at an endpoint and directed inward remains in the system and traverses the interval unless later relabeling changes which physical ant occupies that trajectory.

The motivating classroom scale is

$$L = 2 \text{ feet}, \qquad v = 1 \text{ foot/minute},$$

so the natural time scale is $L/v = 2$ minutes. The analysis below is written for arbitrary L and v.

2 Collision Invariance and Exit Times

Lemma 1 (Collision invariance). *For two identical ants moving at equal speed on a line, collision-and-reversal and pass-through motion produce the same unordered set of positions after the meeting. They differ only in which label is attached to each outgoing trajectory.*

Proof. Immediately before a collision, one ant moves right and the other left. Reversing both velocities at the meeting point gives the same two outgoing geometric paths as allowing both ants to continue straight through the point. The paths are identical as an unlabeled set. □

The local relabeling extends globally without a multiple-collision ambiguity.

Lemma 2 (No higher-order free intersection). *If the initial positions are distinct, no three free trajectories meet at one spacetime point at positive time.*

Proof. Before exit, every free trajectory has the form

$$y_i(t) = x_i + vd_i t, \qquad d_i \in \{-1, +1\}.$$

Among any three trajectories, two have the same slope. If those two met at the same time, equality of their positions would imply equality of their initial positions, contradicting distinctness. □

Any simultaneous pairwise crossings therefore occur at different spatial locations and can be relabeled independently. This makes the pass-through representation exact for unlabeled occupied positions and exit times.

For each ant, define the directed distance to the endpoint it initially faces:

$$\delta_i = \begin{cases} x_i, & d_i = -1, \\ L - x_i, & d_i = +1. \end{cases} \tag{1}$$

Theorem 1 (Exit-time multiset). *Under collision-and-reversal dynamics, the unordered multiset of exit times is*

$$\left\{ \frac{\delta_1}{v}, \dots, \frac{\delta_n}{v} \right\}.$$

If $\delta_{(1)} \leq \cdots \leq \delta_{(n)}$ are the ordered directed distances, then

$$T_{(k)} = \frac{\delta_{(k)}}{v},$$

and the system completion time is

$$T = \frac{1}{v} \max_{1 \leq i \leq n} \delta_i. \tag{2}$$

Proof. The pass-through representation has exactly the same unlabeled evolution as the physical system. Its ith free trajectory reaches the endpoint it faces after δ_i/v, so relabeling preserves the full multiset of exit times. □

Sharp extremal bounds

Since $0 \le \delta_i \le L$,

$$T \le \frac{L}{v}.$$

The bound is sharp. With endpoints allowed, an inward-moving ant at either endpoint has directed distance L, so

$$T_{\max} = \frac{L}{v}.$$

If all initial positions are required to lie in $(0, L)$, L/v remains the supremum but is not attained.

At the other endpoint, the infimum is zero. With distinct open-interval positions, no configuration attains zero, but configurations can approach it arbitrarily closely.

Exercise 1 (Minimum completion time)**.** Assume all ants begin at distinct points in $(0, L)$. Show that for every $\varepsilon > 0$ there is a configuration with $T < \varepsilon$.

Solution. Choose distinct $x_i \in (0, \min\{v\varepsilon, L\})$ and direct every ant left. Then $\delta_i = x_i < v\varepsilon$ for every i, so $T < \varepsilon$. Hence $\inf T = 0$.

For the two-foot, one-foot-per-minute version, the sharp upper bound is

$$\frac{L}{v} = 2 \text{ minutes} = 120 \text{ seconds}.$$

3 Algorithms and Collision Complexity

Equation (2) removes collision handling from the completion-time computation. A single scan computes the maximum directed distance, so completion time is available in $O(n)$ time and $O(1)$ auxiliary space.

Theorem 2 (Optimal completion-time complexity)**.** *In the standard explicit-input model, exact completion time has worst-case complexity* $\Theta(n)$.

Proof. The one-pass maximum gives the upper bound. For the lower bound, suppose an algorithm leaves some ant's input unread. Holding every inspected input fixed, change that ant's directed distance from a small value to a value arbitrarily close to L. The completion time can change while the algorithm's observations do not. Hence every ant must be inspected in the worst case. [1] □

The full ordered exit sequence requires more information. Computing all δ_i/v and comparison-sorting them takes $O(n \log n)$. This is optimal in the comparison model: with $v = 1$ and all ants directed left, arbitrary distinct keys $a_i \in (0, L)$ can be encoded as positions $x_i = a_i$, and the ordered exit times are exactly the sorted keys. Thus the ordered sequence has comparison complexity $\Theta(n \log n)$. [1]

How many collisions can occur?

Sort the initial positions so that $x_1 < \cdots < x_n$. A physical collision occurs exactly when a right-moving free trajectory starting to the left crosses a left-moving free trajectory starting to the right. Therefore

$$C = \#\{(i, j) : i < j,\ d_i = +1,\ d_j = -1\}. \tag{3}$$

Thus C is the inversion count of the binary direction sequence.

If there are r right-moving and $n - r$ left-moving ants, then $C \le r(n - r)$, with equality when all right-moving ants start to the left of all left-moving ants. Hence

$$C_{\max} = \left\lfloor \frac{n^2}{4} \right\rfloor. \tag{4}$$

The physical event sequence can therefore be quadratic even though its completion time is a one-pass statistic.

The random collision count is also explicit.

Theorem 3 (Collision moments under fair directions)**.** *If $D_1, \dots, D_n$ are independent fair signs, then*

$$\mathbb{E}[C] = \frac{n(n-1)}{8}, \tag{5}$$

$$\mathrm{Var}(C) = \frac{n(n-1)(2n+5)}{96}. \tag{6}$$

Consequently,

$$\frac{C}{n^2} \longrightarrow \frac{1}{8}$$

in L^2 and hence in probability.

Proof. Write

$$C = \sum_{i<j} I_{ij}, \qquad I_{ij} = \mathbf{1}_{\{D_i=+1, D_j=-1\}}.$$

Each indicator has mean 1/4 and variance 3/16, giving equation (5). Indicators indexed by disjoint pairs are independent. For each triple $i < j < k$,

$$\mathrm{Cov}(I_{ij}, I_{ik}) = \frac{1}{16}, \qquad \mathrm{Cov}(I_{ij}, I_{jk}) = -\frac{1}{16}, \qquad \mathrm{Cov}(I_{ik}, I_{jk}) = \frac{1}{16}.$$

Their sum is 1/16. Therefore

$$\mathrm{Var}(C) = \binom{n}{2}\frac{3}{16} + 2\binom{n}{3}\frac{1}{16} = \frac{n(n-1)(2n+5)}{96}.$$

After division by n^4, the variance tends to zero while $\mathbb{E}[C]/n^2 \to 1/8$, proving the final claim. □

4 Random Initial Conditions and Reflection Symmetry

The deterministic reduction does not itself require random initial conditions. Randomness enters only through the initial configuration. The most useful structural question is therefore not which direction distribution is convenient, but when changing directions leaves the directed-distance law unchanged.

Let $X_1, \dots, X_n$ be iid on $[0, L]$, and let the direction vector $D = (D_1, \dots, D_n)$ have any probability law on $\{-1, +1\}^n$ that is independent of $X = (X_1, \dots, X_n)$. The directions may be biased and may be mutually dependent.

Theorem 4 (Exact reflection criterion)**.** *The directed exit distances $\Delta_1, \dots, \Delta_n$ are iid with the same law as the initial positions for every joint direction law independent of the positions if and only if the common position law is symmetric about the midpoint:*

$$X_i \overset{d}{=} L - X_i.$$

Proof. For sufficiency, condition on any fixed direction vector $D = d$. Each coordinate Δ_i is then either X_i or $L - X_i$. Midpoint symmetry makes both choices equal in distribution, and deterministic coordinatewise reflections preserve independence because they act separately on the iid coordinates. Thus the conditional joint law of $(\Delta_1, \dots, \Delta_n)$ is the same iid product law for every d, and averaging over the arbitrary law of D leaves it unchanged.

For necessity, choose the deterministic all-left direction vector and the deterministic all-right direction vector. In the first case each directed distance has the law of X_i; in the second it has the law of $L - X_i$. Invariance for every direction law therefore forces $X_i \overset{d}{=} L - X_i$. □

Uniform positions are the canonical special case and give the strongest closed forms:

$$X_1, \dots, X_n \overset{\text{iid}}{\sim} \mathrm{Uniform}(0, L) \quad \Longrightarrow \quad \Delta_1, \dots, \Delta_n \overset{\text{iid}}{\sim} \mathrm{Uniform}(0, L).$$

The fair-independent direction model is therefore sufficient but not necessary for any of the uniform exit-time results below.

Exit order, completion time, and covariance

Let $\Delta_{(1)} \le \cdots \le \Delta_{(n)}$ be the order statistics. Standard uniform order-statistic theory gives[2]

$$\frac{vT_{(k)}}{L} = \frac{\Delta_{(k)}}{L} \sim \mathrm{Beta}(k, n+1-k). \tag{7}$$

Hence

$$\mathbb{E}[T_{(k)}] = \frac{k}{n+1}\frac{L}{v}. \tag{8}$$

For $1 \le i \le j \le n$, the joint covariance structure is

$$\mathrm{Cov}(T_{(i)}, T_{(j)}) = \left(\frac{L}{v}\right)^2 \frac{i(n-j+1)}{(n+1)^2(n+2)}. \tag{9}$$

For $T = T_{(n)}$ and $0 \le t \le L/v$,

$$F_T(t) = \mathbb{P}(T \le t) = \left(\frac{vt}{L}\right)^n, \tag{10}$$

with density

$$f_T(t) = n\left(\frac{v}{L}\right)^n t^{n-1}.$$

For every real $k > 0$,

$$\mathbb{E}[T^k] = \frac{n}{n+k}\left(\frac{L}{v}\right)^k. \tag{11}$$

Therefore

$$\mathbb{E}[T] = \frac{n}{n+1}\frac{L}{v}, \tag{12}$$

$$\mathrm{Var}(T) = \frac{n}{(n+1)^2(n+2)}\left(\frac{L}{v}\right)^2, \tag{13}$$

$$Q_q(T) = \frac{L}{v}q^{1/n}, \qquad 0 < q < 1. \tag{14}$$

For the two-foot classroom scale, $\mathbb{E}[T] = 2n/(n+1)$ minutes; this is $20/11 \approx 1.818$ minutes for $n = 10$ and $100/51 \approx 1.961$ minutes for $n = 50$.

The full spacing law

The individual Beta laws are marginals of a stronger joint result. Put

$$U_{(k)} = \frac{vT_{(k)}}{L}, \qquad U_{(0)} = 0, \qquad U_{(n+1)} = 1,$$

and define the $n+1$ spacings

$$S_k = U_{(k)} - U_{(k-1)}, \qquad k = 1, \dots, n+1.$$

Theorem 5 (Uniform exit-time spacings)**.** *The spacing vector satisfies*

$$(S_1, \dots, S_{n+1}) \sim \mathrm{Dirichlet}(1, \dots, 1). \tag{15}$$

Consequently,

$$\mathbb{E}[S_k] = \frac{1}{n+1}, \qquad \mathrm{Var}(S_k) = \frac{n}{(n+1)^2(n+2)}, \tag{16}$$

and, for $i \ne j$,

$$\mathrm{Cov}(S_i, S_j) = -\frac{1}{(n+1)^2(n+2)}. \tag{17}$$

In physical time, every consecutive expected gap, including the initial and terminal endpoint gaps under the conventions above, is

$$\mathbb{E}\big[T_{(k)} - T_{(k-1)}\big] = \frac{L}{v(n+1)}. \tag{18}$$

Proof. The joint density of the ordered uniforms $(U_{(1)}, \dots, U_{(n)})$ is the constant $n!$ on the simplex

$$0 < u_1 < \cdots < u_n < 1.$$

The change of variables from ordered locations to spacings has unit Jacobian and maps this simplex to

$$s_k \geq 0, \qquad \sum_{k=1}^{n+1} s_k = 1.$$

The resulting constant density is exactly the Dirichlet$(1, \dots, 1)$ density. Its standard moments give equations (16)–(17), and multiplication by L/v gives equation (18). Equation (9) follows as well because each $U_{(k)}$ is a partial sum of the spacings.[2] □

The mean completion-time shortfall from the sharp bound is therefore exactly one expected spacing:

$$\frac{L}{v} - \mathbb{E}[T] = \frac{L}{v(n+1)}. \tag{19}$$

Repeated-trial limit

The expectation in equation (12) answers a configuration-space question. Repeated independent experiments give the corresponding long-run frequency statement.

Theorem 6 (Repeated-trial completion-time limit)**.** *Fix n, L, and v, and let $T_1, T_2, \dots$ be independent completion times generated from the uniform model. Then*

$$\overline{T}_m = \frac{1}{m} \sum_{j=1}^{m} T_j \longrightarrow \frac{n}{n+1} \frac{L}{v} \qquad \textit{almost surely as } m \to \infty. \tag{20}$$

Proof. The completion times are iid and bounded by L/v, so they are integrable. The strong law of large numbers applies, and equation (12) supplies the limiting expectation.[5] □

This limit keeps the number of ants fixed while the number of trials grows. It is distinct from the many-ant limit

$$\lim_{n \to \infty} \mathbb{E}[T] = \frac{L}{v}, \tag{21}$$

whose exact gap is given by equation (19).

The ants-remaining process

At time t, define

$$R(t) = \#\{i : \Delta_i > vt\}.$$

For $0 \leq t \leq L/v$,

$$R(t) \sim \text{Binomial}\left(n, 1 - \frac{vt}{L}\right), \tag{22}$$

so

$$\mathbb{E}[R(t)] = n\left(1 - \frac{vt}{L}\right), \tag{23}$$

$$\text{Var}(R(t)) = n\left(1 - \frac{vt}{L}\right)\frac{vt}{L}. \tag{24}$$

The completion event is exactly $\{R(t) = 0\}$, which recovers equation (10).

The marginal binomial law is only part of the process description.

Proposition 1 (Transition law for ants remaining)**.** *For $0 \leq s < t < L/v$,*

$$R(t) \mid R(s) = r \sim \text{Binomial}\left(r, \frac{L - vt}{L - vs}\right). \tag{25}$$

Thus $R(t)$ is a time-inhomogeneous pure-death Markov process. Its instantaneous transition rate from r to $r-1$ is

$$q_{r,r-1}(t) = \frac{rv}{L - vt}. \tag{26}$$

Proof. Under the uniform model, the individual pass-through exit times are iid Uniform$(0, L/v)$. Conditional on an ant still being present at time s, its exit time is uniform on $(s, L/v)$. Therefore a surviving ant remains present at time t with probability

$$\frac{L/v - t}{L/v - s} = \frac{L - vt}{L - vs}.$$

Conditional on the r survivors at time s, these residual exit times remain independent, giving equation (25). For $h \downarrow 0$, the one-ant conditional exit probability over $(t, t+h]$ is $vh/(L - vt) + o(h)$; multiplying by r gives equation (26). □

Terminal order statistics and extreme values

The usual completion-time limit is the first member of a fuller terminal hierarchy. For a fixed positive integer r, let $T_{(n-r+1)}$ be the rth-to-last exit time.

Theorem 7 (Gamma terminal-gap hierarchy)**.** *For every fixed $r \geq 1$,*

$$n\left(1 - \frac{vT_{(n-r+1)}}{L}\right) \Rightarrow \text{Gamma}(r, 1), \tag{27}$$

where the Gamma distribution is parameterized by shape r and rate 1*.*

Proof. Fix $x \geq 0$, and let $N_n(x)$ count normalized exit times in $(1 - x/n, 1]$. Then

$$N_n(x) \sim \text{Binomial}\left(n, \frac{x}{n}\right) \Rightarrow \text{Poisson}(x).$$

The event

$$n\left(1 - U_{(n-r+1)}\right) > x$$

is equivalent to having at most $r - 1$ observations in $(1 - x/n, 1]$. Hence

$$\mathbb{P}\big(n(1 - U_{(n-r+1)}) > x\big) \longrightarrow e^{-x} \sum_{j=0}^{r-1} \frac{x^j}{j!},$$

which is the survival function of a Gamma$(r, 1)$ random variable. □

For $r = 1$, equation (27) reduces to

$$n\left(1 - \frac{vT}{L}\right) \Rightarrow \text{Exponential}(1), \tag{28}$$

so the gap between completion and the upper endpoint is of order $1/n$. Equivalently, with $M_n = vT/L$ and $Z_n = n(M_n - 1) \leq 0$,

$$\mathbb{P}(Z_n \leq z) \longrightarrow H(z) = \begin{cases} e^z, & z \leq 0, \\ 1, & z > 0. \end{cases}$$

This is a Type III (reverse-Weibull) extreme-value law with generalized-extreme-value shape index $\xi = -1$; the uniform law lies in the Weibull maximum domain of attraction.[3]

5 Beyond Uniform Positions

The reflection theorem isolates the special role of symmetric position laws. Without midpoint symmetry, the direction vector can change the completion-time distribution. The dependence can still be written exactly.

Suppose $X_1, \ldots, X_n$ are iid with continuous CDF F on $[0, L]$, and let D have an arbitrary joint law independent of the positions. Define

$$K = \#\{i : D_i = -1\},$$

the number initially directed left, and for $0 \leq y \leq L$ write

$$a(y) = F(y), \qquad b(y) = 1 - F(L - y).$$

Theorem 8 (General completion-time law)**.** *For $0 \le t \le L/v$,*

$$\mathbb{P}(T \le t) = \mathbb{E}\big[a(vt)^K b(vt)^{n-K}\big] . \tag{29}$$

Consequently,

$$\mathbb{E}[T] = \frac{1}{v}\int_0^L \big\{1 - \mathbb{E}\big[a(y)^K b(y)^{n-K}\big]\big\}\, dy. \tag{30}$$

Proof. Condition on the full direction vector. If $K = k$, then k directed distances have the form X_i and the remaining $n - k$ have the form $L - X_i$. Since the positions are independent,

$$\mathbb{P}(T \le t \mid D) = a(vt)^K b(vt)^{n-K}.$$

Averaging over the direction law gives equation (29). The completion time is nonnegative and bounded by L/v, so applying the tail-integral identity and changing variables $y = vt$ gives equation (30). □

Two useful special cases fall out immediately. If the directions are iid with

$$\mathbb{P}(D_i = -1) = p, \qquad \mathbb{P}(D_i = +1) = 1 - p,$$

then $K \sim \text{Binomial}(n, p)$ and equation (29) becomes

$$\mathbb{P}(T \le t) = G(vt)^n, \qquad G(y) = pF(y) + (1-p)\big[1 - F(L-y)\big]. \tag{31}$$

If instead F is midpoint-symmetric, then $a(y) = b(y)$ and equation (29) reduces to $F(vt)^n$ for *every* direction law independent of the positions, in agreement with the stronger joint reflection-invariance theorem. The uniform model is the case $F(y) = y/L$.

6 Simulation and the Companion Browser Piece

The mathematical reduction has a direct computational use: a simulator does not need to resolve collisions in order to know the exit-time multiset or completion time. The core browser-side calculation is only the directed-distance map:

Listing 1: Exact reduced computation used by the companion implementation.

```
function directedExitTimes(positions, directions, L, v) {
  return positions.map((x, i) =>
    (directions[i] === -1 ? x : L - x) / v
  );
}

function completionTime(positions, directions, L, v) {
  return Math.max(0, ...directedExitTimes(positions, directions, L, v));
}
```

A direct collision animation remains useful because it makes the apparent interaction complexity visible. Its unordered exit times and final completion time should agree with the reduced calculation; disagreement is therefore a diagnostic for collision handling, event ordering, or endpoint logic.

The companion website intentionally combines the two equivalent descriptions from Section 2. The ordinary ants are randomized and rendered on the free pass-through trajectories, so their unlabeled exit behavior is already exact without resolving pairwise reversals. A colored benchmark pair is placed at the two visual ends and explicitly reverses when the pair meets, making the collision rule visible. In the mathematical abstraction the pair represents inward-facing endpoint trajectories; the left benchmark begins at $x = 0$ and traverses the full stick, so the displayed run is forced to the nominal sharp bound L/v regardless of the random ants. The animation is therefore a *hybrid demonstration*, not a sample from the pure iid-uniform completion-time distribution in Section 4. The stochastic formulas and Monte Carlo experiment refer only to the pure random model; the colored pair makes both the collision equivalence and the extremal theorem visible.

The browser animation may use any convenient pixel and wall-clock scale. Equations in the paper depend only on the dimensionless ratios x/L and vt/L, so rescaling the visualization does not change the underlying mathematics. Under the pure random model, equations (22) and (25) give both the marginal and transition laws behind a remaining-ant counter. Under the hybrid website configuration, the two deterministic endpoint trajectories are added to the random component until the terminal time L/v.

The supplied Python and JavaScript source files implement the reduced formulas with explicit input checks. A separate verification script compares the reduction with an event-driven physical simulator on seeded test configurations and exhaustively checks the collision-moment formulas for small n.

7 Monte Carlo and Antithetic Variance Reduction

The closed-form uniform model is useful as a simulation benchmark precisely because the target is known. Normalize

$$U_i = \frac{\Delta_i}{L}, \qquad U_i \overset{\text{iid}}{\sim} \text{Uniform}(0,1),$$

and write

$$M = U_{(n)}, \qquad c = \frac{L}{v}.$$

Then $T = cM$ and

$$\mu = \mathbb{E}[T] = c\frac{n}{n+1}, \qquad \sigma^2 = \text{Var}(T) = c^2\frac{n}{(n+1)^2(n+2)}. \tag{32}$$

For m independent completion times,

$$\widehat{\mu}_m = \frac{1}{m}\sum_{j=1}^{m} T_j$$

is unbiased, with exact variance and standard error

$$\text{Var}(\widehat{\mu}_m) = \frac{n}{m(n+1)^2(n+2)}\left(\frac{L}{v}\right)^2, \tag{33}$$

$$\text{SE}(\widehat{\mu}_m) = \frac{L}{v}\sqrt{\frac{n}{m(n+1)^2(n+2)}}. \tag{34}$$

The iid central limit theorem gives the usual asymptotic normalization.[5]

$$\frac{\sqrt{m}(\widehat{\mu}_m - \mu)}{\sigma} \Rightarrow N(0,1).$$

Exact antithetic calculation

Reflect a generated vector $U = (U_1, \ldots, U_n)$ to

$$U^A = (1-U_1, \ldots, 1-U_n).$$

The reflected completion statistic is

$$M^A = \max_i(1-U_i) = 1 - U_{(1)},$$

so the normalized antithetic pair average is

$$A = \frac{M + M^A}{2} = \frac{U_{(n)} + 1 - U_{(1)}}{2}. \tag{35}$$

Both marginals have the same law as M, so cA is unbiased for μ.

Theorem 9 (Exact antithetic variance)**.** *For the pair average in equation* (35)*,*

$$\text{Corr}(M, M^A) = -\frac{1}{n}, \tag{36}$$

$$\text{Var}(cA) = \left(\frac{L}{v}\right)^2 \frac{n-1}{2(n+1)^2(n+2)}. \tag{37}$$

Proof. For uniform order statistics,[2]

$$\text{Var}(U_{(1)}) = \text{Var}(U_{(n)}) = \frac{n}{(n+1)^2(n+2)},$$

and

$$\text{Cov}(U_{(1)}, U_{(n)}) = \frac{1}{(n+1)^2(n+2)}.$$

Since $M^A = 1 - U_{(1)}$, the covariance of M and M^A is the negative of the latter quantity. Dividing by the common variance gives $-1/n$. Substitution into

$$\mathrm{Var}(A) = \frac{1}{4}\left[\mathrm{Var}(M) + \mathrm{Var}(M^A) + 2\,\mathrm{Cov}(M, M^A)\right]$$

gives equation (37) after multiplication by $(L/v)^2$. □

If m independent base vectors are used, averaging the m pair estimates gives

$$\mathrm{Var}(\widehat{\mu}_m^A) = \frac{1}{m}\left(\frac{L}{v}\right)^2 \frac{n-1}{2(n+1)^2(n+2)}.$$

Against crude Monte Carlo using the same number m of independently generated uniform vectors,

$$\frac{\mathrm{Var}(\widehat{\mu}_m^A)}{\mathrm{Var}(\widehat{\mu}_m)} = \frac{n-1}{2n}. \tag{38}$$

In this problem, M and M^A are obtained in one scan by tracking the maximum and minimum, so no second random vector is required.

A different accounting convention compares one antithetic pair with two independently generated crude functional values. Against $2m$ independent crude realizations,

$$\frac{\mathrm{Var}(\widehat{\mu}_m^A)}{\mathrm{Var}(\widehat{\mu}_{2m})} = \frac{n-1}{n}. \tag{39}$$

For $n = 50$, the two ratios are 0.49 and 0.98. These comparisons answer different cost questions: the first fixes random-vector generation, while the second fixes the number of marginal functional values. Neither is a complete wall-clock model, since random-number generation and min/max updates need not have equal cost. The exact formulas make that accounting choice explicit rather than folding it into the variance-reduction claim. Antithetic variates are standard; here their performance is fully identifiable because the completion-time law is available in closed form.[4]

The accompanying `ants_monte_carlo.py` uses a fixed pseudorandom seed, samples the normalized directed-distance vectors directly, and reports empirical values beside equations (32), (36), (38), and (39).

8 Conclusion

The ants-on-a-stick system is governed by one invariant with consequences at several mathematical scales. Relabeling removes the collision dynamics from every unlabeled exit-time observable, leaving the completion time as the maximum directed distance divided by speed. That same reduction explains why completion time is computable in $\Theta(n)$ despite a worst-case quadratic number of physical collisions, while the full ordered exit sequence retains the $\Theta(n \log n)$ complexity of sorting. Under fair directions, even the collision burden is explicit: its mean and variance are known exactly and C/n^2 concentrates at $1/8$.

Randomization reveals a second structural principle. Midpoint reflection symmetry, rather than fairness or independence of directions, is the essential condition behind direction-law invariance. Uniform positions then identify the complete exit process with uniform order statistics. This gives not only the completion-time law, moments, and quantiles, but the covariance of all ordered exits, a Dirichlet law for inter-exit spacings, an almost-sure repeated-trial limit, an exact pure-death transition law for ants remaining, and a Gamma hierarchy for the final few exit gaps whose first case is the reverse-Weibull completion-time limit. For nonsymmetric position laws, the general formula in equation (29) isolates exactly how the direction-count distribution enters.

The closed-form stochastic law also turns the model into a controlled computational benchmark. Crude and antithetic Monte Carlo estimators can be compared exactly, including the distinction between variance reduction per generated random vector and per marginal functional value. In the companion browser piece, the same mathematics appears visually: randomized ants use the pass-through representation, a colored benchmark pair displays an explicit reversal and realizes the sharp completion-time bound, and the exact exit-process formulas provide the reference model behind the counter and animation. Every extension remains traceable to the original collision invariant, so the recreational problem and the advanced analysis are parts of the same model.

References

[1] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. MIT Press, 4th edition, 2022.

[2] H. A. David and H. N. Nagaraja. *Order Statistics*. Wiley, 3rd edition, 2003.

[3] Paul Embrechts, Claudia Klüppelberg, and Thomas Mikosch. *Modelling Extremal Events: for Insurance and Finance*, volume 33 of *Stochastic Modelling and Applied Probability*. Springer, Berlin, Heidelberg, 1997. ISBN 9783540609315. doi: 10.1007/978-3-642-33483-2.

[4] Paul Glasserman. *Monte Carlo Methods in Financial Engineering*, volume 53 of *Stochastic Modelling and Applied Probability*. Springer, New York, 2004. ISBN 9780387004518. doi: 10.1007/978-0-387-21617-1.

[5] John A. Gubner. *Probability and Random Processes for Electrical and Computer Engineers*. Cambridge University Press, 2006. ISBN 9780521864701.

[6] Harvey Mudd College Department of Mathematics. Ants on a stick. Math Fun Facts, `https://math.hmc.edu/funfacts/ants-on-a-stick/`, n.d. Accessed 3 September 2026.