Journal of Statistical Mechanics:
Theory and Experiment

SISSA

PAPER

Where should you park your car? The $\frac{1}{2}$ rule

To cite this article: P L Krapivsky and S Redner *J. Stat. Mech.* (2020) 073404

View the article online for updates and enhancements.

IOP ebooks™

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection–download the first chapter of every title for free.

This content was downloaded from IP address 129.219.247.33 on 22/02/2021 at 17:39

Journal of Statistical Mechanics: Theory and Experiment
An IOP and SISSA journal

PAPER: Interdisciplinary statistical mechanics

Where should you park your car? The $\frac{1}{2}$ rule

P L Krapivsky[1] **and S Redner**[2]

[1] Department of Physics, Boston University, Boston, MA, 02215 United States of America
[2] Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM, 87501 United States of America

Received 27 March 2020
Accepted for publication 7 May 2020
Published 14 July 2020

Online at stacks.iop.org/JSTAT/2020/073404
https://doi.org/10.1088/1742-5468/ab96b7

Abstract. We investigate parking in a one-dimensional lot, where cars enter at a rate λ and each attempts to park close to a target at the origin. Parked cars also depart at rate 1. An entering driver cannot see beyond the parked cars for more desirable open spots. We analyze a class of strategies in which a driver ignores open spots beyond τL, where τ is a risk threshold and L is the location of the most distant parked car, and attempts to park at the first available spot encountered closer than τL. When all drivers use this strategy, the probability to park at the best available spot is maximal when $\tau = \frac{1}{2}$, and parking at the best available spot occurs with probability $\frac{1}{4}$.

Keywords: traffic and crowd dynamics, stochastic processes, transport properties

Contents

1. Introduction and model 2
2. Spatial distributions of parked cars and vacancies 4
3. Number of vacancies in the active zone 6
4. The best and the actual parking spot 8

© 2020 IOP Publishing Ltd and SISSA Medialab srl
1742-5468/20/073404+14$33.00

5. Cost of parking 11

6. Discussion 12

Acknowledgments 13

References 14

1. Introduction and model

One of the downsides of driving to a popular destination is that good parking spaces near the venue may be hard to find. The dilemma is whether to park far away, which should be easy, and then have a long walk to the destination, or drive close to the venue and then look for a good parking spot, which is likely to be hard. In this work, we investigate a simple class of threshold-driven parking strategies in an idealized parking lot geometry. We find the strategy that maximizes the probability to park in the best available spot.

Because of its pervasive role in our daily lives, understanding parking has been the focus of much study, especially in the urban planning and transportation engineering literatures (see, e.g., [1–7] and references therein). To help in making policy decisions, these investigations typically include many real-world effects, such as practical parking-lot geometries, parking costs, parking limits, and urban planning implications. It is difficult to gain fundamental insights from such reality based approaches. Our approach is to investigate a parsimonious model that captures some key features of parking, and for which some analytic understanding can be gleaned.

For simplicity, the parking lot is one-dimensional, with parking spots labeled by the positive integers $k = 1, 2, \ldots$.and the desired target is located at $k = 0$ (figure 1). Cars enter one at a time from the right at rate λ and each car also departs at rate 1; this input rate is the only parameter of the system. The dynamics of the number of cars is governed by the Poisson process, independent of the parking strategy. However, the spatial distribution of parked cars depends on the parking strategy.

We postulate that a driver cannot see open parking spots closer to the target than the current position of the driver (if there is a contiguous gap of open spots at the current position, the driver can see only to the end of this gap). We additionally assume that after entering the lot, the driver finds a parking space before the next car enters or any of the parked cars depart. While this model is highly idealized, it captures the tradeoff involved in looking for open spots in a crowded parking lot.

The parked cars form groups that are interspersed by gaps (figure 1). Two basic macroscopic observables are the total number of parked cars N and the span L, defined as the distance between the target and the most distant parked car. We are particularly interested in the number of open parking spots within the span (henceforth referred to as vacancies) $V = L - N$, as well as their spatial distribution. While the state of the parking lot is always changing, due to the stochastic arrival and departure of cars, the state of the lot is statistically stationary. Moreover, relative fluctuations become small in the $\lambda \to \infty$ limit. From the underlying Poisson process for the number of parked cars,

J. Stat. Mech. (2020) 073404

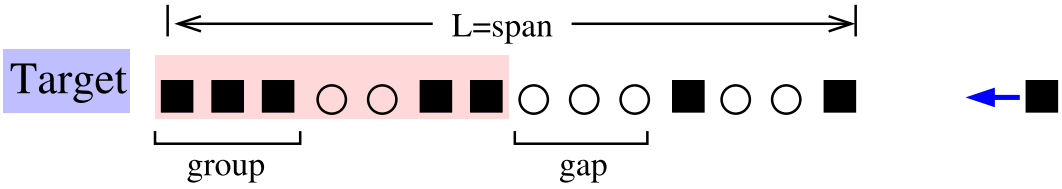

Figure 1. Parking in a one-dimensional lot where cars (squares) enter from the right and circles represent empty spots. Spots beyond the furthest car are not shown. The next car to enter (blue arrow) will start looking for a parking spot in the active zone (red shaded area). In this example, this next entering car parks a distance $k = 4$ from the target, which happens to be the best available spot. Here $N = 7$, $L = 14$, $V = 7$, and $\tau = \frac{1}{2}$.

$\langle N\rangle = \lambda$ and $\langle N^2\rangle - \langle N\rangle^2 = \lambda$, so the relative standard deviation vanishes as $\lambda^{-1/2}$. In real life, parking lots are often nearly full, which corresponds to large λ in our model. Hence we assume $\lambda \gg 1$ unless stated otherwise.

All drivers follow the same parking strategy and we seek the optimal strategy, which we define as the one that maximizes the probability to find the best available spot without backtracking, i.e., the vacancy closest to the target. (If backtracking occurs, the filled vacancy is closest to the target, but there is the added expense of backtracking.) We analyze a class of *threshold* parking strategies that are inspired by threshold strategies that arise in a wide variety of optimal decision problems [8–12]. The classic 'secretary problem' [9–22] is perhaps the best-known example of the utility of a threshold rule. In the framework of our parking problem, the threshold characterizes the degree of risk that an entering driver is willing to accept. We thus define a risk threshold $\tau \in (0, 1)$ as follows: when a car enters the parking lot whose current span is L, the driver ignores all vacancies in $[\tau L, L)$ and only starts looking for vacancies upon reaching a distance τL from the target. The driver parks at the end of the first eligible gap encountered that is closest to the target (figure 1). If no gaps are found when the target is reached, the driver backtracks and parks in the first available vacancy, whose location is necessarily greater than τL. We call the zone $(0, \tau L)$ active since drivers actively look for vacancies there; in the complementary passive zone $[\tau L, L)$, vacancies are initially ignored.

We shall show that in the $\lambda \to \infty$ limit, the probability for a driver to park in the best available spot is maximized by choosing $\tau = \frac{1}{2}$. Normally one does not park at the first vacancy encountered because one feels that better spots will be available. But waiting too long may result in the failure to find an available spot. The $\frac{1}{2}$ Rule provides the best compromise between actually finding a spot and not being a sucker for parking too far away. For the τ-threshold strategy, we will show that the probability to find a single vacancy in the active zone, which is the same as finding the best available parking spot, is $P_1(\tau) = \tau(1 - \tau)$. This probability is maximized when $\tau = \frac{1}{2}$ and the maximum value of $P_1(\tau)$ equals $\frac{1}{4}$.

We previously studied parking in the same one-dimensional geometry [23], in which parking occurred according to either the *optimistic* strategy or to the *prudent* strategy. In the optimistic strategy, the driver goes all the way to the target and then backtracks to the closest available spot. In the prudent strategy, the driver parks at the first gap

J. Stat. Mech. (2020) 073404

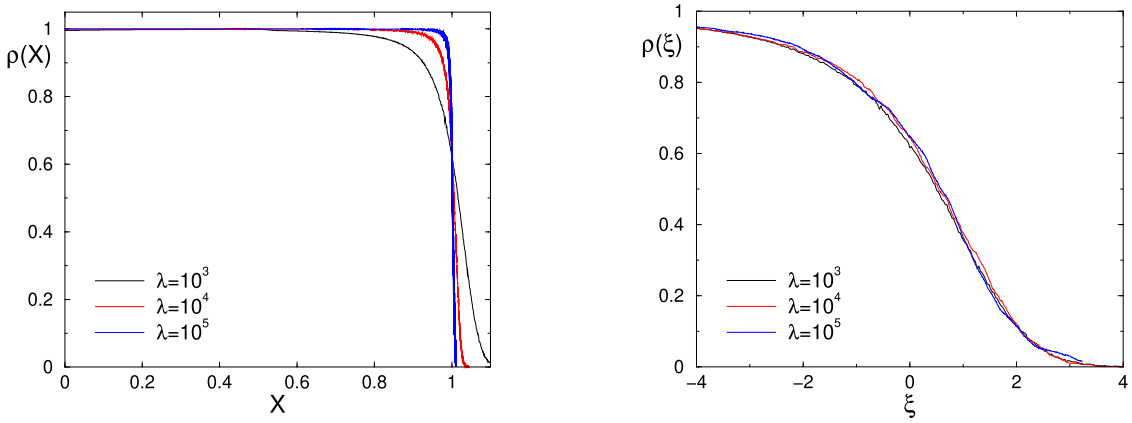

(a) (b)

Figure 2. (a) Simulation data for the density profile of parked cars versus $X = k/\lambda$. (b) The density profile in the front region versus the boundary-layer variable $\xi = (X-1)\lambda^{1/2}$. In (b), the data for $\lambda = 10^4$ and $\lambda = 10^5$ are smoothed over 10-point and 100-point ranges, respectively. For both plots $\tau = \frac{1}{2}$.

encountered; when there are no vacancies, the driver backtracks and parks behind the rightmost parked car. These two rules correspond to the risk threshold $\tau = 0$ and $\tau = 1$, respectively. We will see that these two extremal parking rules are both inferior to the strategy with $\tau = \frac{1}{2}$.

In the next section, we present simulation data for the spatial distribution of cars and the complementary distribution of vacancies. The latter distribution has a rich spatial structure and we offer conjectures for the vacancy density in the active and passive zones that appear to be exact. In section 3 we derive the distribution of the number of vacancies $P_n(\tau)$ in the active zone. In section 4 we argue that the position of the chosen parking spot is spatially uniform, independent of the threshold τ. In section 5 we determine the cost associated with parking and the parking strategy that minimizes this cost. Finally, in section 6 we summarize our results and give some perspectives.

2. Spatial distributions of parked cars and vacancies

For the threshold rule with τ strictly less than 1, the average density of parked cars at position k approaches a step function as the arrival rate $\lambda \to \infty$: $\rho(k) = 1$ for $k < \lambda$ and $\rho(k) = 0$ for $k > \lambda$ (figure 2(a)). That is, the bulk of the parking lot is full and most vacancies are near the far end of the lot where $k \approx \lambda$. Figure 2(b) shows the density profile near the top of the 'Fermi sea' plotted versus the appropriate 'boundary-layer' variable $\xi = (k-\lambda)/\lambda^{1/2}$. The deviation of the spatial distribution from a step function vanishes as $\lambda^{-1/2}$ for $\lambda \to \infty$ and the data for other values of τ are qualitatively similar.

Because the lot is nearly full for large λ, it is more revealing to focus on the spatial distribution of vacancies $1 - \rho(k)$. Since the vacancy density is of the order of $1/\lambda$ in the

J. Stat. Mech. (2020) 073404

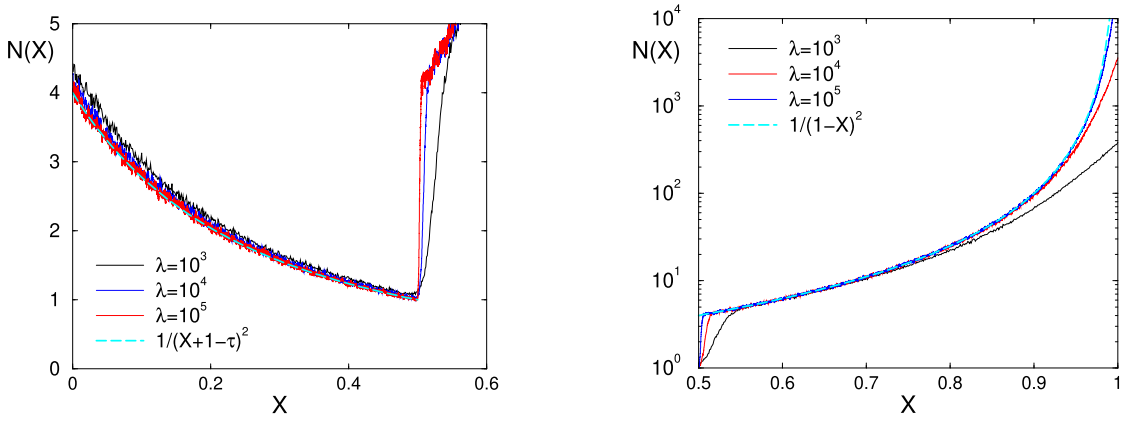

(a) (b)

Figure 3. The scaled density profile of vacancies for $\tau = \frac{1}{2}$ in: (a) the active zone $X < \frac{1}{2}$, and (b) the passive zone $X > \frac{1}{2}$. The data for $\lambda = 10^4$ and $\lambda = 10^5$ are smoothed over a 10-point and 100-point range, respectively.

nearly fully occupied region of the parking lot, it is useful to define the scaled location $X \equiv k/\lambda$ and the scaled vacancy density

$$N(X) = \lambda[1 - \rho(k)]. \tag{1}$$

With this rescaling, we numerically observe that the vacancy density $N_{\rm a}(X)$ in the active zone (defined by $X < \tau$ in scaled units) is well fit by the simple form (figure 3(a)):

$$N_{\rm a}(X) = \frac{1}{(X + 1 - \tau)^2} \quad 0 \leqslant X \leqslant \tau. \tag{2}$$

This conjectured form of the density profile implies that the average number of vacancies in the active zone is

$$\langle n \rangle_{\rm a} = \int_0^\tau {\rm d}X N_{\rm a}(X) = \frac{\tau}{1 - \tau}. \tag{3}$$

This result turns out to be asymptotically exact, and we derive it in the next section without relying on (2).

If there are no vacancies in the active zone, the driver must backtrack and park in the passive zone. Qualitatively, the parking mechanism in the passive zone is the mirror image of parking in the active zone. Based on this insight, as well as on the numerical data itself, we make the guess that the scaled vacancy density in the passive zone $\tau < X < 1$ is given by

$$N_{\rm p}(X) = \frac{1}{(1 - X)^2} \quad \tau < X < 1. \tag{4}$$

This simple form also provides an excellent fit to the data in the spatial portion of the parking lot that is nearly full (figure 3(b)).

J. Stat. Mech. (2020) 073404

The jump in the vacancy density profile at $X = \tau$ reflects the nature of the threshold strategy. Since a driver starts looking for a vacancy only when $X = \tau$ is reached, parking spots just closer than $X = \tau$ are likely to be taken, while the best spots close to $X = 0$ are likely to be 'wasted'. Conversely, if one needs to backtrack to park, it is the spots just beyond $X = \tau$ that will be taken, while more distant parking spots will remain more plentiful. A curious and unexplained feature of the density profiles (2) and (4) is that $N_{\rm a}(X = 0) = N_{\rm p}(X = \tau)$.

From the density (4), we estimate the number of vacancies in the passive zone by

$$\langle n\rangle_{\rm p} = \int_{\tau}^{1-\epsilon} {\rm d}X\, N_{\rm p}(X). \tag{5}$$

Here, we introduce an upper cutoff to the integral because the vacancy density (4) cannot hold all the way to $X = 1$. Since the density profile has a transition zone whose width is of the order of $\lambda^{1/2}$, the cutoff in scaled units should be $\epsilon \sim \lambda^{-1/2}$. With this cutoff, (5) gives $\langle n\rangle_{\rm p} \sim \lambda^{1/2}$. Thus for large λ and with drivers all following the same threshold strategy, there are of the order of $\lambda^{1/2}$ lousy parking spots more distant than the threshold, and a few good parking spots within the threshold.

3. Number of vacancies in the active zone

We now derive the distribution of the number of vacancies in the active zone, from which we deduce the optimality of the $\frac{1}{2}$ rule. Denote by $P_n(\tau)$ the probability to find n vacancies in the active zone of the parking lot, $0 < X < \tau$, when the threshold is τ. Further, define $V_1(X_1\,;\tau)$ as the probability density for a single vacancy at X_1 and generally $V_n(X_1, \ldots, X_n\,;\tau)$ as the probability density for n vacancies at $0 < X_1 < X_2 < \cdots < X_n < \tau$. The probability to find n vacancies in the active zone is therefore

$$P_n = \int \cdots \int_{0<X_1<\cdots<X_n<\tau} {\rm d}X_1 \ldots {\rm d}X_n V_n(X_1, \ldots, X_n\,;\tau). \tag{6}$$

Hereinafter we typically drop the dependence on τ to declutter the notation. Explicitly

$$P_1 = \int_0^{\tau} {\rm d}X_1 V_1(X_1), \tag{7a}$$

$$P_2 = \int_0^{\tau} {\rm d}X_1 \int_{X_1}^{\tau} {\rm d}X_2 V_2(X_1, X_2), \tag{7b}$$

$$P_3 = \int_0^{\tau} {\rm d}X_1 \int_{X_1}^{\tau} {\rm d}X_2 \int_{X_2}^{\tau} {\rm d}X_3 V_3(X_1, X_2, X_3), \tag{7c}$$

etc. For notational consistency, we also write $V_0 = P_0$ for the probability of no vacancies in the active zone.

To compute the probability distribution P_n for the number of vacancies in the active zone, we ostensibly need the densities V_n. As we now show, our approach circumvents

J. Stat. Mech. (2020) 073404

the need for the complete information about these densities. The quantity V_0 satisfies the rate equation

$$\frac{\mathrm{d}V_0}{\mathrm{d}t} = -\lambda\tau V_0 + \lambda\int_0^\tau \mathrm{d}X_1 V_1(X_1). \tag{8}$$

The loss term on the right accounts for the departure of a car from the fully occupied active zone. The gain term accounts for a car that parks in the single vacancy within the active zone. In the steady state, (8) becomes

$$\int_0^\tau \mathrm{d}X_1 V_1(X_1) = \tau V_0. \tag{9}$$

The left-hand side of (9) equals P_1 (see (7a)), so we have

$$P_1 = \tau V_0. \tag{10}$$

Similarly to (8), we write the rate equation for $V_1(X_1)$. In the steady state, this equation reduces to

$$\int_{X_1}^\tau \mathrm{d}X_2 V_2(X_1, X_2) + V_0 = \tau V_1(X_1) + V_1(X_1). \tag{11}$$

The first term on the left side accounts for parking when two vacancies exist, while the second term accounts for the departure of a car from the fully occupied active zone. The first term on the right accounts for the departure of a car when the active zone contains a single vacancy, while the second term accounts for a car that parks at X_1. We now integrate (11) over the active zone, $0 < X_1 < \tau$. Using (7b), the left-hand side becomes $P_2 + \tau V_0$, while the right-hand side becomes $(1+\tau)P_1$, after recalling (7a). Hence $P_2 = (1+\tau)P_1 - \tau V_0$, which, in conjunction with (10), yields

$$P_2 = \tau^2 V_0. \tag{12}$$

Continuing this line of reasoning, we write the rate equation for $V_2(X_1, X_2)$. In the steady state, this equation reduces to

$$(\tau+1)V_2(X_1, X_2) = \int_{X_2}^\tau \mathrm{d}X_3 V_3(X_1, X_2, X_3) + V_1(X_1) + V_1(X_2). \tag{13}$$

Integrating (13) over $0 < X_1 < X_2 < \tau$ ultimately yields

$$P_3 = \tau^3 V_0. \tag{14}$$

Following this approach for general n, we find $P_n = \tau^n V_0$. The normalization condition, $\sum_{n\geqslant 0} P_n = 1$, fixes V_0 to be $1-\tau$. Thus the distribution of the number of vacancies in the active zone is the geometric distribution

$$P_n(\tau) = (1-\tau)\tau^n. \tag{15}$$

The average number of vacancies in the active zone, $0 < X < \tau$, is therefore

$$\langle n\rangle_{\mathrm{a}} = \sum_{n\geqslant 0} n P_n(\tau) = \frac{\tau}{1-\tau}. \tag{16}$$

J. Stat. Mech. (2020) 073404

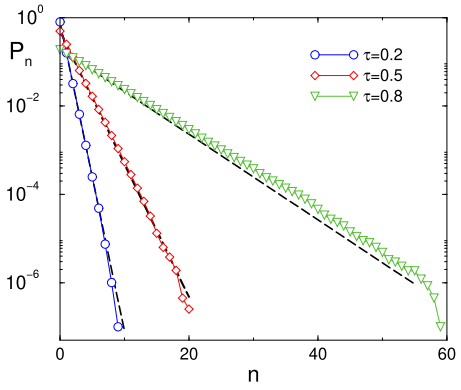

Figure 4. The probability P_n for n vacancies in the active zone for three threshold values and for $\lambda = 10^4$. The dashed lines are the theoretical predictions from (15).

Our simulation data in figure 4 is in excellent agreement with the prediction (15).

From (15), the probability to have one vacancy in the active zone is $P_1(\tau) = \tau(1-\tau)$. By its definition, this quantity coincides with the probability that a newly entering car parks in the best available spot. This probability is maximized when $\tau = \tau_* = \frac{1}{2}$. At this optimal threshold value, there is one vacancy, on average, in the active zone and the probability of actually parking in the best available spot is $P_1(\tau_*) = \frac{1}{4}$.

4. The best and the actual parking spot

A natural question that arises upon entering a crowded parking lot is: where is the best open parking spot? Let $B(X)$ be the probability density to have the best parking spot at X. This probability density has a non-trivial spatial structure (figure 5). For $\tau = 0$ (the optimistic strategy in reference [23]), the distribution is flat: $B(X) = 1$. That is, if one drives straight to the target and then starts looking for a parking spot by backtracking, the closest parking spot is equiprobably anywhere in the lot.

For the general situation, $0 < \tau < 1$, the distribution $B(X)$ exhibits different behaviors in the active and passive zones, with a jump at $X = \tau$. In the passive zone, which gets occupied due to backtracking, the distribution is again flat; moreover, $B_{\rm p}(X) = 1$ for $X > \tau$. In the active zone, $X < \tau$, the probability density $B_{\rm a}(X)$ is a decreasing function of X which becomes progressively more peaked as τ increases. That is, if one starts looking for parking spots when X reaches τ, it is the parking spots with $X \lesssim \tau$ that are more likely to be filled, while the desirable spots near $X = 0$ will be relatively unfilled. The probability density to have the best spot next to the target, $B_{\rm a}(0)$, is obviously equal to the vacancy density $N_{\rm a}(0)$; this implies that $B_{\rm a}(0) = (1-\tau)^{-2}$. The probability density $B(X)$ is normalized, so that $\int_0^\tau {\rm d}X B_{\rm a}(X) = \tau$. These are two exact properties of the probability density $B_{\rm a}(X)$ in the active zone. The challenge is to determine $B_{\rm a}(X)$ analytically.

The threshold-$\frac{1}{2}$ rule leads to success, viz, to parking in the best available spot, in a fraction $P_1(\frac{1}{2}) = \frac{1}{4}$ of attempts; otherwise parking occurs in a suboptimal spot.

J. Stat. Mech. (2020) 073404

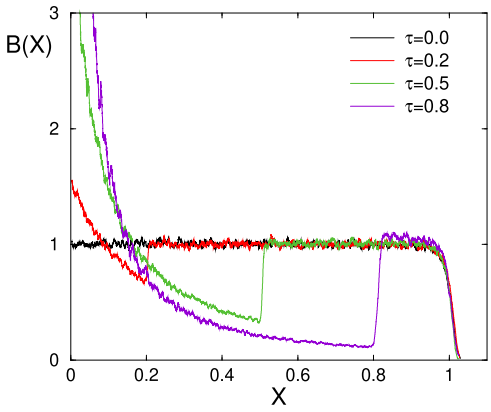

Figure 5. The probability density $B(X)$ that the best parking spot is at a scaled distance X from the target for threshold strategies with $\tau = 0, 0.2, 0.5, 0.8$. The data are based on $\lambda = 10^4$ and are smoothed over a 50-point range.

Thus a useful measure of the efficacy of the threshold strategy is the ratio $r \equiv X/X_{\rm B}$ of the *actual* parking location X to the best available location $X_{\rm B}$. Denote by $q(r)$ the probability density that $X/X_{\rm B} = r$. We know that $r = 1$ if there is either 1 or 0 vacancies in the active zone. In these two cases, the car parks in the best available spot without backtracking or with backtracking, respectively. The two events occur with probability $P_0 + P_1 = 1 - \tau^2$. Hence

$$q(r) = (1 - \tau^2)\delta(r - 1) + Q(r). \tag{17}$$

The continuous part $Q(r)$ of the probability density accounts for the situations when there is more than one vacancy in the active zone, so the entering car necessarily parks in a suboptimal spot and $r > 1$.

The continuous part $Q(r)$ of the probability density satisfies $\int_1^\infty {\rm d}r\, Q(r) = \tau^2$ and it appears to have a power-law tail for large r (figure 6) in the general situation when $0 < \tau < 1$. For the threshold values shown in this figure, the data is reasonably fit by the power law $r^{-\nu}$, with $\nu \approx 1.8$. That is, even though a driver typically does not park in the best spot, it is likely that parking occurs close to the best spot. When $\tau = 0$ (the optimistic strategy in [23]), the car necessarily parks in the best spot by backtracking: $q(r) = \delta(r - 1)$.

Although the ratio of the actual parking location to the best location has a non-trivial distribution, the spatial distribution of the actual parking location $\Pi(X)$ is remarkably simple: $\Pi(X) = 1$ (figure 7). To understand this striking result, consider the part of this distribution $\Pi_{\rm a}(X)$ in the active zone, $X \leqslant \tau$, when a car finds a parking spot without backtracking. The strategy in the active zone is effectively the same as the strategy with an effective threshold $\tau_{\rm eff} = 1$ in the case where the span is τL. The properly scaled spatial variable is X/τ and hence

$$\Pi_{\rm a}(X) = \Phi(y), \quad y = \frac{X}{\tau}. \tag{18a}$$

J. Stat. Mech. (2020) 073404

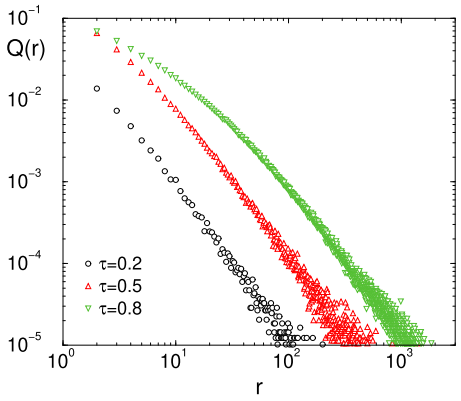

Figure 6. The continuous part $Q(r)$ of the probability density of the ratio of the actual parking location to the best parking location, $r = X/X_{\rm B}$.

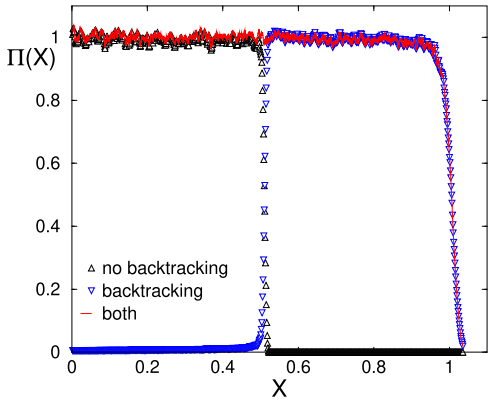

Figure 7. The spatial distribution of parking locations for $\tau = \frac{1}{2}$. The data have been smoothed over a 50-point range and only every tenth data point is shown.

When backtracking occurs, cars enter the passive zone and park in the first gap, so effectively they follow the prudent strategy with span $(1-\tau)L$. The properly scaled spatial variable is now $(1-X)/(1-\tau)$ as the cars move to right when they backtrack. Thus the part of the parking distribution $\Pi_{\rm p}(X)$ that corresponds to the passive zone is

$$\Pi_{\rm p}(X) = \Phi(y), \quad y = \frac{1-X}{1-\tau}. \tag{18b}$$

The fractions of cars that park in the active and passive zones are τ and $1-\tau$. These fractions also equal $\int_0^\tau {\rm d}X \Pi_{\rm a}(X)$ and $\int_\tau^1 {\rm d}X \Pi_{\rm p}(X)$, respectively. Using (18a) and (18b) we indeed recover τ and $1-\tau$ after imposing the normalization condition

$$\int_0^1 {\rm d}y\, \Phi(y) = 1. \tag{19}$$

J. Stat. Mech. (2020) 073404

To determine the form of the probability density $\Pi(X)$, let us compute its derivatives at the threshold location $X = \tau$:

$$\left.\frac{\mathrm{d}^n\Pi}{\mathrm{d}X^n}\right|_{X=\tau-0} = \tau^{-n}\Phi^{(n)}(1),$$

$$\left.\frac{\mathrm{d}^n\Pi}{\mathrm{d}X^n}\right|_{X=\tau+0} = (-1)^n(1-\tau)^{-n}\Phi^{(n)}(1),$$

where $\Phi^{(n)}(y) = \mathrm{d}^n\Phi(y)/\mathrm{d}y^n$. The derivatives from left and write coincide only if $\Phi^{(n)}(1) = 0$ for all $n \geqslant 1$. Thus $\Pi(X)$ is smooth at $X = \tau$ when all derivatives of $\Phi(y)$ vanish at $y = 1$. The vanishing of all derivatives implies[3] that the distribution is uniform, and the actual value $\Phi(y) = 1$ is fixed by the normalization condition (19). Our above argument is not a proof, but it shows that *postulating* analyticity of the probability density $\Pi(X)$ at the threshold $X = \tau$ suffices to derive its uniformity.

5. Cost of parking

Maximizing the probability to park in the best available spot is natural and may be compatible with the intrinsic irrationality of human behavior in parking lots see, e.g., [24]). However, the threshold parking rule is not necessarily the most rational. A more logical approach is to minimize the cost of parking, which we define as the arrival time to reach the target. This arrival time is the sum of the time spent in driving in the lot to find the parking spot plus the walking time from the parking spot to the target. In appropriate units where the walking speed equals 1, the walking time is just the distance from the parking spot to the target, while the driving time is the driving distance in the lot divided by the driving speed. This driving time also equals the driving distance multiplied by the ratio of the walking speed to the driving speed in the lot; we denote this ratio by ϵ.

For the τ-threshold parking strategy, these arrival times are

$$\frac{\tau}{2} + \epsilon\left(1 - \frac{\tau}{2}\right) \quad \text{and} \quad \frac{1+\tau}{2} + \epsilon\left(1 + \frac{1+\tau}{2}\right),$$

for the cases of parking in the active and passive zones, respectively. Here we have made use of the result that the parking location is uniformly distributed in both the active and passive zones. Multiplying these arrival times by the probabilities for these two types of parking events, τ and $1-\tau$ respectively, the average parking cost is

$$\frac{1}{2}(1+3\epsilon) - \epsilon\tau^2. \tag{20}$$

This cost function attains a minimum with the prudent strategy, $\tau = 1$. Despite being optimal according to the above choice for the cost function, the prudent strategy is

[3] We tacitly assume that $\Phi(y)$ is analytic, namely its Taylor series (centered at $y = 1$) converges. There are infinitely differentiable functions which are not analytic; for such function, vanishing of all derivatives does not assure that the function is constant.

J. Stat. Mech. (2020) 073404

psychologically upsetting to most people. In the prudent strategy, the parking spot that a driver takes is essentially never the best available, as the typical number of vacancies is of order $\sqrt{\lambda}$. Perhaps even more upsetting to a driver who has just parked is that the absolutely best parking spot—the one that is adjacent to the target—will be available in approximately 89% of all realizations [23].

In deriving (20), we tacitly assumed that $\epsilon > 0$; however, the limiting case of $\epsilon = 0$ is actually quite natural. Setting $\epsilon = 0$ is equivalent to only counting the time spent in walking to the target in the cost of parking. Since moving in a car requires little physical effort, drivers with limited walking ability, as well as drivers with little children, may find this cost function more suitable. Furthermore, if the target is a store were the drivers tend to buy heavy goods, it is again reasonable to minimize the time of walking from the target to the parked car, i.e., set $\epsilon = 0$. In this case, the average parking cost is independent of strategy. Since the $\frac{1}{2}$ rule maximizes the probability of snagging the best parking spot, which everyone prefers, it is rational to adopt the threshold-$\frac{1}{2}$ rule if the cost of walking through the parking lot is prohibitively high.

6. Discussion

We investigated a class of threshold parking strategies in which a driver who enters a parking lot ignores all open spots a distance greater than τL from the target. Starting at a distance τL, which defines the end of the active zone, the driver parks at near end of the first gap encountered. Here L is the distance from the target to the last parked car and τ is the risk threshold. If there are no available spots in the desirable active zone, the driver backtracks and parks in the first spot encountered by backtracking.

When the ratio of the arrival rate to departure rate of each car is large, $\lambda \gg 1$, open parking spots are rare in the active zone, and the spots that do exist are more likely to be close to the target (figure 2(b)). Our primary result is that the probability for a newly entering car to park at the best available spot is maximized for $\tau = \frac{1}{2}$. At this optimal threshold, the probability of parking in this best spot is $\frac{1}{4}$. We conjecture that our strategy is the best amongst all (deterministic), and not necessarily threshold parking strategies. Proving or refuting this conjecture represents an appealing challenge.

Our analysis for number of vacancies in the active zone (section 3) tacitly assumed that $0 < \tau < 1$; however, our main result (15) can be continued to $\tau = 0$ and $\tau = 1$; these limits correspond to the optimistic and prudent strategies in our earlier study [23]. This continuation correctly predicts that, for the optimistic strategy, the probability of finding the best spot without backtracking vanishes. Indeed, for this strategy, backtracking occurs with probability $\lambda/(1+\lambda)$ [23], so backtracking is asymptotically certain in this case. For the prudent strategy, the probability of backtracking vanishes as λ^{-1}, but the probability that the spot where the car actually parks is closest to the target vanishes as $\lambda^{-1/2}$. Thus the $\frac{1}{2}$ rule is superior to both the optimistic and prudent strategies in finding the best available parking spot.

While we found an optimal parking strategy, we were unable to analytically determine several natural spatial properties of the parked cars, such as the distribution of the number of vacancies V, the number of gaps (figure 1), the distribution of the

J. Stat. Mech. (2020) 073404

span L, and multisite densities, such as $V_n(X_1, \ldots, X_n)$ in the active zone. Additionally, the expressions (2) and (4) for the vacancy densities in the active and passive zones, which appear to be correct, are conjectural. The decoupling approximation, which we employed in [23] to determine the average density of cars (or vacancies) at a given spatial location, fails for the threshold strategy and another analytical approach is needed.

Our threshold parking rules are inspired by the strategies that arise in optimal stopping problems and in decision theory (see, e.g., [8–12]). Some of the methods developed in this body of work, particularly in the realm of the famous 'secretary problem' [9–22], may prove useful to analyze optimal parking. Indeed, the best strategy for various decision theory optimization problems is often a threshold strategy [8, 13].

An intriguing contribution to the work on optimal stopping problems is the odds algorithm [17], which is summarized by adage: sum the odds to one and stop. The odds algorithm has been proved for some optimal stopping problems with independent events [17–19]. We now show that the odds algorithm holds in our case where parking events are not independent. By definition, the odds are the ratios $p/(1-p)$ where the p's are the probabilities of success. In our situation, the probability of successful parking at spot k is the probability $1-\rho(k)$ that this spot is empty[4]. Since the $\rho(k)$ are very close to 1 in the bulk of the span, the odds coincide with the probabilities (i.e., the vacancy densities) to leading order. Summing the vacancy densities in the active zone yields the average number of vacancies, $\langle n\rangle_{\rm a} = \tau/(1-\tau)$. Equating this number to one recovers the $\frac{1}{2}$ rule.

It is worth pointing out that our model assumes a homogeneous population of drivers; namely, they all have the same parking threshold parameter τ. Rich behaviors may emerge for heterogeneous populations. The simplest heterogeneous population consists of drivers that each have independent random parking thresholds that, for example, are uniformly distributed on $(0, 1)$. There is no longer an optimization problem to be solved, but the spatial properties of the parked cars could be interesting and tractable.

This heterogeneous model is a sort of equilibrium counterpart of the hashing problem that was originally introduced by Konheim and Weiss [25] and described by Knuth [26] using the parking problem language. This hashing problem and its numerous extensions have been subsequently studied by many authors, see, e.g., [27–31] and references therein. In the standard application of the hashing problem, one attempts to write a fixed-size file onto a computer disk. A random location on the disk is picked and if the file can be accommodate there, the file is written. If not, the next location is selected. If this location is empty, the file is written there. If not, continue moving in one direction until a vacancy is encountered. The list of starting locations is the hash table. In the language of parking, the lot is finite, there is no departure mechanism (files are not erased), so eventually the parking lot becomes full. Richer behaviors arise when the file sizes are variable. The correspondence between hashing and parking offers the possibility that the substantial research literature on the hashing problem could guide developments in our parking problem.

[4] When a driver meets the first gap in the active zone, parking occurs at the spot in the gap that is closest to the target; more distant vacancies in the gap are ignored. However, gaps in the active zone typically consists of isolated vacancies. Indeed, the total number of gaps with ℓ vacancies scales as $\lambda^{-(\ell-1)}$ and hence vanishes in the $\lambda\to\infty$ limit for all $\ell \geqslant 2$.

J. Stat. Mech. (2020) 073404

Acknowledgments

PLK thanks the hospitality of the Santa Fe Institute where this work was initiated. SR gratefully acknowledges financial support from NSF Grant DMR-1608211. We also thank John Miller for helpful conversations.

References

[1] Van Der Groot D 1982 A model to describe the choice of parking places *Transp. Res.* A **16** 109
[2] Young W, Thompson R G and Taylor M A P 1991 A review of urban car parking models *Transp. Rev.* **11** 63
[3] Axhausen K W and Polak J W 1991 Choice of parking: stated preference approach *Transportation* **18** 59
[4] Thompson R G and Richardson A J 1998 A parking search model *Transp. Res.* A **32** 159
[5] Arnott R and Rowse J 1999 Modeling parking *J. Urban Econ.* **45** 97
[6] Teodorović D and Lučić P 2006 Intelligent parking systems *Eur. J. Oper. Res.* **175** 1666
[7] Klappenecker A, Lee H and Welch J L 2014 Finding available parking spaces made easy *Ad Hoc Netw.* **12** 243
[8] Lindley D V 1961 Dynamic programming and decision theory *J. R. Stat. Soc.* **10** 39
[9] DeGroot M H 1970 *Optimal Statistical Decisions* (New York: McGraw-Hill)
[10] Chow Y S, Robbins H and Siegmund D 1971 *Great Expectations: The Theory of Optimal Stopping* (Boston: Houghton Mifflin Company)
[11] Berezovsky B A and Gnedin A V 1984 *Problems of Best Choice* (Moscow: Akademia Nauk)
[12] Preater J 1999 The best-choice problem for partially ordered objects *Oper. Res. Lett.* **25** 187
[13] Dynkin E 1963 Optimal choice of the instant for stopping a Markov process *Sov. Math. Dokl.* **150** 238–40
[14] Chow Y S, Moriguti S, Robbins H and Samuels S M 1964 Optimal selection based on relative rank (the 'secretary problem') *Israel J. Math.* **2** 81
[15] Gilbert J P and Mosteller F 1966 Recognizing the maximum of a sequence *J. Am. Statist. Assoc.* **61** 35
[16] Ferguson T S 1989 Who solved the secretary problem? *Statist. Sci.* **4** 282
[17] Bruss F T 2000 Sum the odds to one and stop *Ann. Probab.* **28** 1384
[18] Bruss F T 2003 A note on bounds for the odds theorem of optimal stopping *Ann. Probab.* **31** 1859
[19] Dendievel R 2013 New developments of the odds theorem *Math. Sci.* **38** 111
[20] Georgiou N, Kuchta M, Morayne M and Niemiec J 2008 On a universal best choice algorithm for partially ordered sets *Random Struct. Algorithm* **32** 263
[21] Freij R and Wästlund J 2010 Partially ordered secretaries *Electron. Commun. Probab.* **15** 504
[22] Garrod B and Morris R 2013 The secretary problem on an unknown poset *Random Struct. Algorithm* **43** 429
[23] Krapivsky P L and Redner S 2019 Simple parking strategies *J. Stat. Mech.* 093404
[24] Vanderbilt T 2008 *Traffic: Why We Drive the Way We Do (And What it Says about Us)* (New York: Knopf)
[25] Konheim A G and Weiss B 1966 An occupancy discipline and applications *SIAM J. Appl. Math.* **14** 1266
[26] Knuth D E 1973 *The Art of Computer Programming* (Sorting and Searching vol 3) (Reading, MA: Addison-Wesley)
[27] Flajolet P, Poblete P and Viola A 1998 On the analysis of linear probing hashing *Algorithmica* **22** 490
[28] Janson S 2001 Asymptotic distribution for the cost of linear probing hashing *Random Struct. Algorithm* **19** 438
[29] Janson S 2005 Individual displacements for linear probing hashing with different insertion policies *ACM Trans. Algorithm* **1** 177
[30] Flajolet P and Sedgewick R 2009 *Analytic Combinatorics* (Cambridge: Cambridge University Press)
[31] Lackner M-L and Panholzer A 2016 Parking functions for mappings *J. Comb. Theory* A **142** 1

J. Stat. Mech. (2020) 073404