

QUANTUM GRAVITY AND INVENTORY ACCUMULATION

BY SCOTT SHEFFIELD

Massachusetts Institute of Technology

We begin by studying inventory accumulation at a LIFO (last-in-first-out) retailer with two products. In the simplest version, the following occur with equal probability at each time step: first product ordered, first product produced, second product ordered, second product produced. The inventory thus evolves as a simple random walk on \mathbb{Z}^2 . In more interesting versions, a p fraction of customers orders the “freshest available” product regardless of type. We show that the corresponding random walks scale to Brownian motions with diffusion matrices depending on p .

We then turn our attention to the *critical Fortuin–Kastelyn random planar map model*, which gives, for each $q > 0$, a probability measure on random (discretized) two-dimensional surfaces decorated by loops, related to the q -state Potts model. A longstanding open problem is to show that as the discretization gets finer, the surfaces converge in law to a limiting (loop-decorated) random surface. The limit is expected to be a *Liouville quantum gravity* surface decorated by a *conformal loop ensemble*, with parameters depending on q . Thanks to a bijection between decorated planar maps and inventory trajectories (closely related to bijections of Bernardi and Mullin), our results about the latter imply convergence of the former in a particular topology. A phase transition occurs at $p = 1/2$, $q = 4$.

1. Introduction. A planar map is a connected planar graph together with an embedding into the sphere (which we identify with the complex plane $\mathbb{C} \cup \{\infty\}$), defined up to topological deformation. Self loops and multi-edges are allowed. See Figure 1.

There is a significant literature on *random* planar maps of various types. As we illustrate later, planar maps are in one-to-one correspondence with planar quadrangulations, which in turn can be interpreted as Riemannian surfaces, obtained by gluing together unit squares along their boundaries (see Figure 6). If we also declare a random subset of the edges of a planar map to be “open”, we can generate a set of loops on the Riemannian surface that separate open clusters from dual clusters (see Figures 7 and 9). We are interested in the “scaling limits” of these random loop-decorated surfaces when the number of unit squares tends to infinity.

However, we approach the problem from an unconventional direction. We begin in Sections 2 and 3 by stating and proving a theorem about inventory accumulation

Received June 2014; revised September 2015.

MSC2010 subject classifications. 60K35, 60D05.

Key words and phrases. Liouville quantum gravity, planar map, scaling limit, Schramm–Loewner evolution, FK random cluster model, continuum random tree, mating of trees.




FIG. 1. A planar map embedded in the plane.

at a LIFO (last-in-first-out) retailer with two products. Given a certain model for production and sales, in which a p fraction of customers always “orders fresh” (i.e., takes the most recently added product regardless of type), we find that the time evolution of the two product inventories scales to a two-dimensional Brownian motion with a diffusion matrix depending on p . (See Theorem 2.5 and Figure 3.) We view this result as interesting in its own right. We also find a surprising phase transition: when $p \geq 1/2$, there are no macroscopic inventory *imbalances*. That is, while the total inventory fluctuates by about \sqrt{n} after n steps, the *difference* between the two product inventories fluctuates by $o(\sqrt{n})$.

Section 4 then presents a bijection between inventory accumulation trajectories and possible instances of so-called critical Fortuin–Kasteleyn (FK) random planar maps. These are random planar maps—together with random distinguished edge subsets—whose laws depend on a parameter $q \in (0, \infty)$ (closely related to the q -state Potts model), which turns out to be related to p by $p = \frac{\sqrt{q}}{2+\sqrt{q}}$. As mentioned above, we may interpret these maps as loop-decorated surfaces. From this point of view, the inventory-trajectory central limit theorem mentioned above will become a scaling limit theorem about the corresponding loop-decorated surfaces. Specifically, it will describe (an infinite volume version of) the limiting law of a certain tree and dual tree that are canonically associated to a discretized loop-decorated surface and that encode the structure of the surface. Our bijection specializes to a classical bijection of Mullin when the distinguished edge subset is required to form a spanning tree of the planar map [6, 47]. It is equivalent to a construction by Bernardi when the structure of the planar map is fixed (see Section 4 of [8]).¹

¹Precisely, Bernardi constructs a many-to-one map from the set of planar maps with distinguished edge subsets to the set of planar maps with distinguished spanning trees and provides a way to understand the pre-image of a *single* planar map with a distinguished spanning tree. Mullin’s bijection gives a way of enumerating the set of *all* planar maps with distinguished spanning trees. At the end

The proofs in this paper are discrete and elementary, relying only on standard results in probability (such as the optional stopping theorem). We remark, however, that we view this paper as part of a larger program to relate random planar maps to continuum objects such as Liouville quantum gravity and the Schramm–Loewner evolution. We will not discuss these ideas in the body of the paper, but we provide a brief explanation of this point of view in an [Appendix](#). Let us note that since this paper was first posted to the arXiv, there have been many additional papers that have used the theory developed here to prove the convergence the random planar map models to continuum models (involving conformal loop ensembles and Liouville quantum gravity) in certain topologies [19, 27–30, 46]. Other works have described tail exponents of the planar map models [5], a generalization to more than two burger types [42], a further study of the infinite critical-FK random map [12], and a related sandpile result [55].

We also remark that there is a vast operations research literature on inventory management protocols, including various schemes involving two products, but to our knowledge this is the first paper to address this particular model, and also the first to make the connection between inventory trajectories and planar maps.

2. Inventory trajectories: Setup and theorem statement. In this section, we describe a random walk on a particular semi-group (related, as we will see in Section 4, to random planar maps) and study its scaling limit. We interpret the walk as a simple model for inventory accumulation at a LIFO (last-in-first-out) retailer with two product types and three order types (first product, second product, and “flexible”/“freshest available”) arriving at random times. As a convenient mnemonic, we refer to the two products as *hamburgers* and *cheeseburgers* (or *burgers* collectively).²

2.1. Defining the semigroup. Write $\Theta = \{\textcircled{C}, \textcircled{H}, \boxed{C}, \boxed{H}, \boxed{F}\}$. We view Θ as an alphabet of symbols that represent, respectively, a cheeseburger, a hamburger, a cheeseburger order, a hamburger order, and a “flexible” (either hamburger or cheeseburger) order. Informally, a word W in Θ describes a day in the life of a restaurant: for example, the word $\textcircled{C}\textcircled{H}\boxed{C}\boxed{C}\boxed{F}\boxed{C}\textcircled{C}\boxed{H}\textcircled{H}$ above describes

of [7], Bernardi suggests (as a future research project) trying to combine his construction with one of the available enumerations of spanning-tree-decorated planar maps, such as Mullin’s bijection. This is essentially what we do: our inventory trajectory encoding can be understood as a particularly simple way to unify the Mullin and Bernardi constructions.

²We use this particular metaphor in part for linguistic simplicity (one noun, two prefixes) and in part because the chute stacked with hamburgers and cheeseburgers (produced in a back room, ordered in a front room) is something of an inventory management icon.

a day in which first someone produced a cheeseburger (to put on the top of a “stack” of burgers), then someone produced a hamburger, then two people ordered cheeseburgers, one ordered “freshest available,” another ordered a cheeseburger, someone produced a cheeseburger, someone ordered a hamburger, and someone produced a hamburger. Informally, we say that whenever a burger is produced, it is added to the top of the stack, and whenever an order is placed it is fulfilled (if possible) by removing the highest matching burger from the stack (but an order that cannot be filled immediately remains unfilled—because the impatient customer leaves the restaurant, say).

To describe this formally, we view the elements of Θ as generators for a certain (associative) semigroup \mathcal{G} , each element of which can be represented by a word W in the alphabet Θ (with the empty string \emptyset as a left and right identity). This is the semigroup of words modulo four “order fulfillment” relations

$$\boxed{C}\boxed{C} = \boxed{H}\boxed{H} = \boxed{C}\boxed{F} = \boxed{H}\boxed{F} = \emptyset,$$

and two “commutativity” relations

$$\begin{aligned} \boxed{C}\boxed{H} &= \boxed{H}\boxed{C}, \\ \boxed{H}\boxed{C} &= \boxed{C}\boxed{H}. \end{aligned}$$

A word W in Θ is called *reduced* if no element of $\{\boxed{C}, \boxed{H}\}$ appears to the left of any element of $\{\boxed{C}, \boxed{H}, \boxed{F}\}$. In other words, it contains a list of orders followed by a list of burgers, for example,

$$W = \boxed{C}\boxed{H}\boxed{H}\boxed{F}\boxed{H}\boxed{F}\boxed{H}\boxed{C}\boxed{H}\boxed{H}.$$

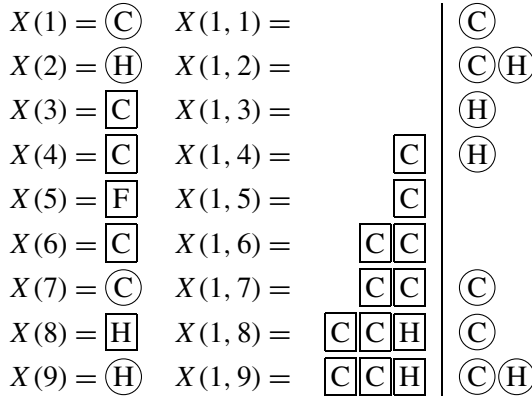
PROPOSITION 2.1. *For each finite-length word W in Θ there is a unique reduced word \overline{W} that is equivalent to W modulo the above relations. Thus, \mathcal{G} is in one-to-one correspondence with the set of reduced words.*

PROOF. Before establishing the uniqueness part of Proposition 2.1, we establish existence by noting that for each $W \in \Theta$ we can explicitly construct a specific reduced word \overline{W} inductively as follows. If W has length zero or one, then we take $\overline{W} = W$. Now suppose \overline{W} has been defined for all W of length k or less. Then we extend the definition to words of length $k + 1$ as follows. Write $\overline{W}\boxed{C} = \overline{W}\boxed{C}$ and $\overline{W}\boxed{H} = \overline{W}\boxed{H}$, and let $\overline{W}\boxed{C}$, $\overline{W}\boxed{H}$ or $\overline{W}\boxed{F}$ be obtained by excising (respectively) the rightmost \boxed{C} , \boxed{H} , or “either \boxed{C} or \boxed{H} ” symbol from \overline{W} —or, if no such symbols exists, appending (respectively) a \boxed{C} , \boxed{H} or \boxed{F} to the right of the list of orders in \overline{W} .

To give a concrete example, consider a map $X : \mathbb{Z} \rightarrow \Theta$. For $a \leq b$ we will write

$$X(a, b) := \overline{X(a)X(a+1)X(a+2) \cdots X(b)}.$$

The following illustrates the sequences $X(1, k)$ corresponding to a particular $X(k)$ sequence:



If we interpret $W = X(1)X(2) \cdots X(n)$ as a “day in the life of a restaurant,” then the corresponding $\overline{W} = X(1, n)$ (in this case $\boxed{C} \boxed{C} \boxed{H} \textcircled{C} \textcircled{H}$) contains the unfulfilled orders (in the order they were added—again, we assume that customers leave the restaurant without a burger if their orders cannot be fulfilled immediately) followed by the unconsumed burgers (in the order they were added). One may imagine the unconsumed burgers to be arranged in a “stack” or a “chute”.³

Next, we claim that if W_1 and W_2 are equivalent words (modulo the order fulfillment and commutativity relations (i.e., they correspond to the same element of \mathcal{G})) then we will have $\overline{W}_1 = \overline{W}_2$. To see this, observe that inserting a consecutive pair of the form $\textcircled{C} \boxed{C}$, $\textcircled{H} \boxed{H}$, $\textcircled{C} \boxed{F}$ or $\textcircled{H} \boxed{F}$ into a word W has no effect on \overline{W} , and also that reversing the order of an adjacent pair of the form $\textcircled{C} \textcircled{H}$ or $\textcircled{H} \textcircled{C}$ in W has no effect on \overline{W} . This (and the fact that $\overline{W} = W$ when W is reduced) establishes the uniqueness in Proposition 2.1 and implies that \mathcal{G} is in one-to-one correspondence with the set of reduced words. \square

We also note that the procedure used to construct the reduced form of $X(1, n)$ inductively from X has a “reverse direction” analog. Consider the following ex-

³An Internet search for “burger chute” turns up many images of multi-lane chutes (which allow burgers to be sorted by type). We stress that our model envisions a *single-lane* chute in which each burger is added at the top (so that the order of the remaining burgers corresponds to the order in which they were added) and each order is filled from the top (by taking away the highest suitable burger and letting the remaining burgers slide down).

ample

$$\begin{array}{lcl}
 X(0) = \boxed{C} & X(0, 0) = & \boxed{C} \\
 X(-1) = \textcircled{H} & X(-1, 0) = & \boxed{C} \textcircled{H} \\
 X(-2) = \textcircled{H} & X(-2, 0) = & \boxed{C} \textcircled{H} \textcircled{H} \\
 X(-3) = \textcircled{C} & X(-3, 0) = & \textcircled{H} \textcircled{H} \\
 X(-4) = \boxed{F} & X(-4, 0) = & \boxed{F} \textcircled{H} \textcircled{H} \\
 X(-5) = \boxed{C} & X(-5, 0) = & \boxed{C} \boxed{F} \textcircled{H} \textcircled{H} \\
 X(-6) = \textcircled{H} & X(-6, 0) = & \textcircled{H} \textcircled{H} \\
 X(-7) = \textcircled{C} & X(-7, 0) = & \textcircled{H} \textcircled{H} \\
 X(-8) = \textcircled{C} & X(-8, 0) = & \textcircled{C} \textcircled{H} \textcircled{H}
 \end{array}$$

From this point of view, when we add a new $X(-k) \in \{\boxed{C}, \textcircled{H}, \boxed{F}\}$, we are adding this element to the left of the stack of orders, and when we add a new $X(-k) \in \{\textcircled{C}, \textcircled{H}\}$, we are annihilating the leftmost matching order (or adding the burger to the left of the list of unconsumed burgers, if there is no available matching order). The sequence $X(1, k)$ has a “burger stack” whose length can go up and down as k increases but an “order stack” whose length can only increase as k increases; similarly, the sequence $X(-k, 0)$ has an order stack whose length can go up and down as k increases but a burger stack whose length can only increase as k increases.

2.2. *Defining matches in random infinite words.* Let $X(n)$ be i.i.d. random variables, indexed by $n \in \mathbb{Z}$, each of which takes values in $\{\textcircled{C}, \textcircled{H}, \boxed{C}, \textcircled{H}, \boxed{F}\}$ with respective probabilities $\{1/4, 1/4, (1 - p)/4, (1 - p)/4, p/2\}$ for some fixed $p \in [0, 1]$ (so that a p fraction of the orders are of type \boxed{F}). Let μ denote the corresponding probability measure on the space Ω of maps from \mathbb{Z} to Θ .

If a burger that is added at time m is consumed at time n , we will say that m and n are a *match* and write $m = \phi(n)$ and $n = \phi(m)$. In other words, if we consider $X(m') \cdots X(n')$ to represent a “day in the life of the restaurant” as described above for any $m' \leq m$ and $n' \geq n$, then m and n are a match if and only if the burger added at time m is consumed at time n on that day. This definition does not depend on the values of $X(k)$ for $k < m$ or $k > n$, since it is equivalent to the statement that $X(m) \in \{\textcircled{C}, \textcircled{H}\}$ and n is the smallest integer greater than m for which $X(m, n)$ does not contain a burger of the same type as $X(m)$. This in turn holds if and only if one of the following four things happens:

1. $X(m) = \textcircled{H}$, $X(n) = \boxed{H}$, and $X(m + 1, n - 1)$ is a word containing only the letters \textcircled{C} and \boxed{C} , for example,

$$X(m + 1, n - 1) = \boxed{C} \boxed{C} \boxed{C} \boxed{C} \textcircled{C} \textcircled{C} \textcircled{C} \textcircled{C} \textcircled{C}.$$

2. $X(m) = \textcircled{C}$, $X(n) = \boxed{C}$, and $X(m + 1, n - 1)$ is a word containing only the letters \textcircled{H} and \boxed{H} , for example,

$$X(m + 1, n - 1) = \boxed{H}\boxed{H}\boxed{H}\boxed{H}\textcircled{H}\textcircled{H}\textcircled{H}.$$

3. $X(m) = \textcircled{H}$, $X(n) = \boxed{F}$ and $X(m + 1, n - 1)$ is a word containing only the letter \boxed{C} , for example,

$$X(m + 1, n - 1) = \boxed{C}\boxed{C}\boxed{C}\boxed{C}.$$

4. $X(m) = \textcircled{C}$, $X(n) = \boxed{F}$ and $X(m + 1, n - 1)$ is a word containing only the letter \boxed{H} , for example,

$$X(m + 1, n - 1) = \boxed{H}\boxed{H}\boxed{H}\boxed{H}\boxed{H}\boxed{H}\boxed{H}\boxed{H}.$$

If $X(m) \in \{\textcircled{C}, \textcircled{H}\}$ and m has no match [i.e., the burger added at time m is never consumed—which would be the case, for example, if we had $X(k) \in \{\textcircled{C}, \textcircled{H}\}$ for all $k > m$]—we write $\phi(m) = \infty$. If $X(n) \in \{\boxed{C}, \boxed{H}, \boxed{F}\}$ and n has no match (i.e., the order at time n is unfulfilled, no matter how far back in time one starts), then we write $\phi(n) = -\infty$.

For example, in the sequence

$X(1) = \textcircled{C}$	$X(1, 1) =$	\textcircled{C}
$X(2) = \textcircled{H}$	$X(1, 2) =$	$\textcircled{C}\textcircled{H}$
$X(3) = \boxed{C}$	$X(1, 3) =$	\textcircled{H}
$X(4) = \boxed{C}$	$X(1, 4) =$	$\boxed{C}\textcircled{H}$
$X(5) = \boxed{F}$	$X(1, 5) =$	\boxed{C}
$X(6) = \boxed{C}$	$X(1, 6) =$	$\boxed{C}\boxed{C}$
$X(7) = \textcircled{C}$	$X(1, 7) =$	$\boxed{C}\boxed{C}\textcircled{C}$
$X(8) = \boxed{H}$	$X(1, 8) =$	$\boxed{C}\boxed{C}\textcircled{H}\textcircled{C}$
$X(9) = \textcircled{H}$	$X(1, 9) =$	$\boxed{C}\boxed{C}\textcircled{H}\textcircled{C}\textcircled{H}$

described above, we have $\phi(3) = 1$; $\phi(1) = 3$ and $\phi(2) = 5$; $\phi(5) = 2$, but the values of $\phi(4), \phi(6), \phi(7), \phi(8), \phi(9)$ necessarily lie outside the interval $\{1, 2, \dots, 9\}$ and are not determined by $X(1), X(2), \dots, X(9)$.

PROPOSITION 2.2. *It is μ almost surely the case that for every $m \in \mathbb{Z}$, we have $\phi(m) \notin \{-\infty, \infty\}$. In other words, every $X(j)$ has a unique match, almost surely, so that ϕ is an involution on \mathbb{Z} .*

PROOF. We first claim that this holds whenever $X(m) \in \{\textcircled{C}, \textcircled{H}\}$. Observe that the net number of burgers (i.e., the number of burger symbols minus the number of order symbols) added between times m and n , as a function of n , is a simple

random walk on \mathbb{Z} . It follows that there will almost surely exist values of $n > m$ for which this quantity is arbitrarily negative, and hence $X(m + 1, n)$ contains an arbitrarily long sequence of orders. [Recall that the number of orders in $X(m + 1, n)$ is nondecreasing as n increases.] If $X(m) = \textcircled{C}$, then the first time that a \boxed{C} or \boxed{F} is added to this list of orders will be a time at which $X(m)$ is consumed; thus, on the event that m has no match it is almost surely the case that the number of \boxed{H} orders in $X(m + 1, n)$ tends to infinity as $n \rightarrow \infty$ while the number of \boxed{C} orders remains zero. If this happens for some m , then there cannot be an m' for which the same thing happens with the roles of hamburgers and cheeseburgers reversed [since $X(m + 1, m')$ is a fixed finite length word—appending it on the left cannot remove an arbitrarily long sequence of \boxed{C} elements from the corresponding reduced word]. Thus, it is almost surely the case that either every cheeseburger added (at any integer time) is ultimately consumed or every hamburger added (at any integer time) is ultimately consumed. Since each of these two events is translation invariant, the zero-one law for translation invariant events (see any textbook with an introduction to ergodic decompositions, e.g., [25]) implies that each has probability zero or probability one. As observed above, the union of these two events has probability one, so by symmetry each of them separately has μ probability one. Thus, μ a.s. every burger of either type is ultimately consumed, which implies the claim. A similar argument (in the reverse direction) shows that $\phi(n)$ is a.s. finite whenever $X(n) \in \{\boxed{C}, \boxed{H}, \boxed{F}\}$. \square

The following is an immediate consequence of the above construction.

PROPOSITION 2.3. *The reduced word $X(1, n)$ contains precisely those $X(k)$ corresponding to the $k \in \{1, 2, \dots, n\}$ for which $\phi(k) \notin \{1, 2, \dots, n\}$. In other words, it contains the list of unmatched orders [in order of appearance in the sequence $X(1), X(2), \dots, X(n)$] followed by the ordered list of unmatched burgers [in order of appearance in $X(1), X(2), \dots, X(n)$].*

2.3. Infinite stacks and random walks: Main theorem. By analogy with Proposition 2.3, we define $X(-\infty, n)$ to be the ordered sequence of $X(k)$ for which $k \leq n$ but $\phi(k) > n$. Informally, we can write

$$X(-\infty, n) = \overline{\cdots X(-3)X(-2)X(-1)X(0)X(1)X(2) \cdots X(n)},$$

and interpret $X(-\infty, n)$ as the reduced form word corresponding to the product of all $X(k)$ with $k \leq n$. We view $X(-\infty, n)$ as a semi-infinite stack of \textcircled{C} and \textcircled{H} symbols, indexed by the negative integers; it has a unique top element but no bottom element. [The number of elements in $X(-\infty, n)$ is μ a.s. infinite for all n by Proposition 2.2—if it were some finite value k , then with positive probability there would be a consecutive sequence of $k + 1$ orders after time n , and at least one order would have no match.]

We similarly write

$$X(n, \infty) := \overline{X(n)X(n+1), \dots},$$

which we interpret to mean the ordered sequence of $X(k)$ for which $k \geq n$ but $\phi(k) < n$. It is natural to represent $X(n, \infty)$ as a sequence of $\{\boxed{C}, \boxed{H}, \boxed{F}\}$ values indexed by the positive integers. While the $X(-\infty, n)$ can be interpreted as a semi-infinite stack of burgers waiting to be consumed (as n increases in time), the stack $X(n, \infty)$ can be interpreted as a semi-infinite queue of customers waiting to be served (as n decreases in time). A useful equivalent definition is that $X(n, \infty)$ is the limit of the order stacks of $X(n, m)$ (which only increase in length as m increases, with new orders being added on the right), as $m \rightarrow \infty$. Similarly, $X(-\infty, n)$ is the limit of the burger stacks of $X(m, n)$ (which only increase in length as m decreases, with new burgers being added on the left) as $m \rightarrow -\infty$.

Next, define

$$Y(n) := \begin{cases} X(n), & X(n) \in \{\textcircled{C}, \textcircled{H}, \boxed{C}, \boxed{H}\}, \\ \boxed{C}, & X(n) = \boxed{F}, X(\phi(n)) = \textcircled{C}, \\ \boxed{H}, & X(n) = \boxed{F}, X(\phi(n)) = \textcircled{H}. \end{cases}$$

In other words, $Y(n)$ is obtained from $X(n)$ by replacing each \boxed{F} with a \boxed{C} or \boxed{H} , depending on which burger type was actually consumed by the \boxed{F} order. For $a \leq b$, we also write

$$Y(a, b) := \overline{Y(a)Y(a+1) \cdots Y(b)},$$

and observe that this is the same as $X(a, b)$ except that each \boxed{F} is replaced with the corresponding \boxed{C} or \boxed{H} symbol.

For every word W in the symbols $\{\textcircled{C}, \textcircled{H}, \boxed{C}, \boxed{H}, \boxed{F}\}$, we write $\mathcal{C}(W)$ for the net burger count (i.e., the number of $\{\textcircled{C}, \textcircled{H}\}$ symbols minus the number of $\{\boxed{C}, \boxed{H}, \boxed{F}\}$ symbols in W). Analogously, if W has no \boxed{F} symbols, then we define $\mathcal{D}(W)$ to be the net discrepancy of hamburgers over cheeseburgers (i.e., the number of $\{\textcircled{H}, \boxed{C}\}$ symbols minus the number of $\{\textcircled{C}, \boxed{H}\}$ symbols).

DEFINITION 2.4. Given the infinite $X(n)$ sequence, let \mathcal{C}_n be the integer valued process defined by $\mathcal{C}_0 = 0$ and $\mathcal{C}_n - \mathcal{C}_{n-1} = \mathcal{C}(Y(n))$ for all n . Similarly, write $\mathcal{D}_0 = 0$ and $\mathcal{D}_n - \mathcal{D}_{n-1} = \mathcal{D}(Y(n))$. Thus, \mathcal{C}_n and \mathcal{D}_n keep track of the net change in the burger count and the burger discrepancy since time zero. When $n \geq 0$ we have $\mathcal{C}_n = \mathcal{C}(Y(1, n))$ and $\mathcal{D}_n = \mathcal{D}(Y(1, n))$. [For this purpose, we will write $Y(1, 0) = \emptyset$ by convention.] When $n < 0$, we have $\mathcal{C}_n = -\mathcal{C}(Y(n+1, 0))$ and $\mathcal{D}_n = -\mathcal{D}(Y(n+1, 0))$. As a shorthand and slight abuse of notation, we will also write, when a and b are integers,

$$\begin{aligned} \mathcal{C}(a) &= \mathcal{C}(Y(a)), & \mathcal{C}(a, b) &= \mathcal{C}(Y(a, b)), \\ \mathcal{D}(a) &= \mathcal{D}(Y(a)), & \mathcal{D}(a, b) &= \mathcal{D}(Y(a, b)). \end{aligned}$$




FIG. 2. Above: top few elements on the stack $X(-\infty, n - 1)$ represented by an up-left edge for each cheeseburger and an up-right edge for each hamburger. Below: top few elements of $X(-\infty, n)$, depending on value of $X(n)$. The uppermost dot above (shown slightly larger) is positioned at a location $A_{n-1} = (\mathcal{D}_{n-1}, \mathcal{C}_{n-1})$ describing the net change in inventory since time zero. The new uppermost dots below are positioned at locations $A_n = (\mathcal{D}_n, \mathcal{C}_n)$.

We also write $A_n = (\mathcal{D}_n, \mathcal{C}_n)$ for integer n . We extend the definition to the real numbers (by piecewise linear interpolation) so that $t \rightarrow A_t$ is an infinite continuous path.

Figure 2 suggests a natural way to visualize the evolution of $X(-\infty, n)$: here $X(-\infty, n)$ is represented by a stack of up-left edges (one for each \textcircled{C}) and up-right edges (one for each \textcircled{H}) ending at the location A_n . If $p = 1$, then every time \mathcal{C}_n decreases, it backtracks along $X(-\infty, n)$, and every time \mathcal{C}_n increases, the corresponding change in \mathcal{D}_n is decided by a coin toss. (This is related to a process called the *Brownian snake*; see, e.g., [36].) In this case, for each fixed n , the burger stack $X(-\infty, n)$ has the law of a sequence of i.i.d. elements of $\{\textcircled{C}, \textcircled{H}\}$. When $p = 0$, the process A_n is just a simple random walk on the set of integer pairs (x, y) for which $x + y$ is even, and the stack $X(-\infty, n)$ encodes information about the past history of A_n . (Informally, the “head of the snake” moves around and drags the body along with it in a way that maintains the requirement that the body always consists of a sequence of up-left and up-right edges, as in Figure 2.) When $p \in (0, 1)$ the transition probabilities are averages of the extreme cases $p = 0$ and $p = 1$.

Our main scaling limit result concerns the random process A_t of Definition 2.4.

THEOREM 2.5. *As $\varepsilon \rightarrow 0$, the random functions $\varepsilon A_{t/\varepsilon^2}$ converge in law (with respect to the L^∞ metric on any compact interval) to*

$$(1) \quad (B_{\alpha t}^1, B_t^2),$$

where B_t^1 and B_t^2 are independent standard one-dimensional Brownian motions and

$$\alpha = \max\{1 - 2p, 0\}.$$

In particular, the above theorem implies that C_n scales to ordinary Brownian motion, which is not surprising since it is just a simple random walk on \mathbb{Z} regardless of p . When $p = 0$, the processes \mathcal{D}_n and C_n are independent simple random walks on \mathbb{Z} , and it is also unsurprising that $A_n = (\mathcal{D}_n, C_n)$ scales to an ordinary two-dimensional Brownian motion in this case. When $p = 1$ (and one only removes burgers from the top of the stack, because all orders are flexible), the law of $X(-\infty, n)$ (for any fixed n) is that of an i.i.d. sequence of \textcircled{C} and \textcircled{H} values. In this case, the fact that inventory changes tend to be “well balanced” between hamburgers and cheeseburgers [so that the first term in (1) is identically zero and $\varepsilon A_{t/\varepsilon^2}$ concentrates on the vertical axis as $\varepsilon \rightarrow 0$] can be deduced from the law of large numbers. Indeed, when $p = 1$, the magnitude of C_n has order $n^{1/2}$ with high probability, while the magnitude of \mathcal{D}_n has order $n^{1/4} = o(n^{1/2})$. To see where the $n^{1/4}$ comes from, recall that $C_0 = 0$ and condition on $m := \min\{C_j : 0 \leq j \leq n\}$ and on C_n , noting that $C_0 - m$ and $C_n - m$ are both of order \sqrt{n} . Note that the top $C_0 - m$ burgers in the time zero stack fill the orders remaining in the reduced word $X(1, n)$, and the top $C_n - m$ burgers in the time n stack are the types of the burgers in $X(1, n)$. The types of the top $C_0 - m$ burgers in the time zero stack and the top $C_n - m$ burgers in the time n stack can then be determined with independent coin tosses, so that \mathcal{D}_n is of order $\sqrt{\sqrt{n}} = n^{1/4} = o(n^{1/2})$.

The theorem states that as long as $p \geq 1/2$ this balance continues to hold: when n is large, the net inventory accumulation between time zero and n is close to evenly divided between hamburgers and cheeseburgers, with high probability. When $p < 1/2$ the fluctuations of \mathcal{D}_n are on the same order as those of C_n . Put differently, a LIFO retailer accumulates major inventory discrepancies (on the same order as the total inventory fluctuation) if and only if more than half of its customers have a product preference. Figure 3 illustrates sample trajectories of A_n for both $p < 1/2$ and $p > 1/2$.

3. Inventory trajectories: Constructions and proofs. The primary goal of this section is to prove Theorem 2.5. On the way, we will establish some independently interesting lemmas involving properties of “excursion words” of various types, equivalent formulations of the $p = 1/2$ phase transition, typical lengths of random reduced words, and monotonicity properties for the time evolution of an inventory stack.

We recall that the *optional stopping theorem* states that if X_0, X_1, X_2, \dots is a martingale (resp., supermartingale) and K is a bounded stopping time then $\mathbb{E}[X_K] = \mathbb{E}[X_0]$ (resp., $\mathbb{E}[X_K] \leq \mathbb{E}[X_0]$). Furthermore, if K is an arbitrary (possibly unbounded) stopping time and X_1, X_2, \dots is a martingale (or supermartingale) whose values are bounded below then $\mathbb{E}[X_K] \leq \mathbb{E}[X_0]$. By symmetry, this also implies that if K is an arbitrary stopping time and X_1, X_2, \dots is a martingale (or submartingale) whose values are bounded above then $\mathbb{E}[X_K] \geq \mathbb{E}[X_0]$. We will use these facts frequently.




FIG. 3. Sample trajectory traced by $A_n = (\mathcal{D}_n, \mathcal{C}_n)$ as n ranges from 0 to 100,000 when $p = 0.25$ (left). Similar trajectories (axes not shown) when $p = 0.75$ (middle) and $p = 1$ (right). In each case we use the same initial data: $X(-\infty, 0)$ is an infinite stack alternating between burger types. In each case $X(1), \dots, X(100,000)$ are then chosen from μ and $(\mathcal{D}_n, \mathcal{C}_n)$ updated accordingly for $n \in \{1, 2, \dots, 100,000\}$. As n tends to infinity, and both axes are scaled by $1/\sqrt{n}$, the left trajectory tends to a two dimensional Brownian motion with variance $1 - 2p = 0.5$ times as large in the left-right direction as the up-down direction. The middle and right trajectories tend to one-dimensional Brownian motions, concentrated on a vertical axis. (See Remark 3.17.)

TABLE 1
Relationship between p , χ and a variance limit

$p \leq 1/2$	$\chi = 2$	$\lim_{n \rightarrow \infty} n^{-1} \text{Var}[\mathcal{D}_n] = 1 - 2p$
$p \geq 1/2$	$\chi = 1/p$	$\lim_{n \rightarrow \infty} n^{-1} \text{Var}[\mathcal{D}_n] = 0$

3.1. *A variance calculation.* Let J be the smallest positive integer for which the (reduced) word $X(-J, -1)$ has a nonempty burger stack (i.e., at least one \textcircled{C} or \textcircled{H} symbol) and let $\chi = \chi(p)$ be the expected length of $X(-J, -1)$. As preparation for proving Theorem 2.5, we will prove the following lemma.

LEMMA 3.1. *If $p \leq 1/2$, then $\chi = 2$ and $\text{Var}[\mathcal{D}_n] = (1 - 2p)n + o(n)$, while if $p \geq 1/2$ then $\chi = 1/p$ and $\text{Var}[\mathcal{D}_n] = o(n)$, as summarized in Table 1.*

This section will prove Lemma 3.1 modulo one other lemma (Lemma 3.2) to be established later. The rough idea behind the argument is pretty simple. On the event that $X(0) = \boxed{F}$, it is not hard to see that $\phi(0) = -J$, as defined above. Our intuition suggests that the addition of extra \boxed{F} symbols to the model should decrease the fluctuation of inventory imbalances. Basically, this is because on the event that a fresh order at time 0 consumes a cheeseburger [i.e., $X(0) = 0$ and $X(-J) = \textcircled{C}$], one might guess that, on average, there had been some net surplus of \textcircled{C} during the period immediately before time 0. In particular, one would expect the word $X(-J, -1)$ to contain a net surplus on \textcircled{C} symbols. Quantifying this effect precisely will amount to computing the expectation of $\mathcal{D}(-J, -1)$ on the event $\phi(-J) = \textcircled{C}$, which will turn out (see 1 in the list below) to be equivalent to computing $\chi = \mathbb{E}[|X(-J, -1)|]$. The arguments below will follow this idea and will give the reader a quick idea of where the $\alpha = \max\{1 - 2p, 0\}$ of Theorem 2.5 comes from and why a phase transition occurs at $p = 1/2$. We present them as a series of observations:

1. The order stack of $X(-J, -1)$ must consist of zero or more orders of type opposite to the one burger in $X(-J, -1)$, for example, $X(-J, -1) = \boxed{C}\boxed{C}\boxed{C}\textcircled{H}$ or $\boxed{H}\boxed{H}\textcircled{C}$ or \textcircled{C} . In particular, the length of the word $X(-J, -1)$ is given by

$$(2) \quad |X(-J, -1)| = |\mathcal{D}(-J, -1)| = -\mathcal{C}(-J, -1) + 2.$$

2. $\mathcal{C}(-j, -1)$ is a martingale in j satisfying $\mathcal{C}(-j, -1) \leq 1$ for all $j \in \{1, 2, \dots, J\}$ and $\mathbb{E}[\mathcal{C}(-1, -1)] = 0$. By the optional stopping theorem for martingales bounded above, the expectation of a martingale at a (possibly unbounded) stopping time is at least the expectation at time zero (and at most the upper bound). This gives $\mathbb{E}[\mathcal{C}(-J, -1)] \in [0, 1]$, and hence by (2),

$$(3) \quad \chi := \mathbb{E}[|X(-J, -1)|] = \mathbb{E}[|\mathcal{D}(-J, -1)|] \in [1, 2].$$

3. By (2) and (3), we have $\mathbb{E}[\mathcal{C}(-J, -1)] = \mathbb{E}[\mathcal{C}(-1, -1)] = 0$ if and only if $\chi = 2$. Since the optional stopping theorem implies

$$\mathbb{E}[\mathcal{C}(-1, -1)] = \mathbb{E}[\mathcal{C}(-J, -1)1_{J \leq n}] + \mathbb{E}[\mathcal{C}(-n, -1)1_{J > n}],$$

and the former term on the RHS tends to $\mathbb{E}[\mathcal{C}(-J, -1)]$ we have

$$(4) \quad \chi = 2 \quad \text{if and only if} \quad \lim_{n \rightarrow \infty} \mathbb{E}[\mathcal{C}(-n, -1)1_{J > n}] = 0.$$

4. By (3), the expectation $\mathbb{E}[\mathcal{D}(-J, -1)]$ exists and, by symmetry,

$$(5) \quad \mathbb{E}[\mathcal{D}(-J, -1)] = 0.$$

By (4) and the fact that $|\mathcal{D}(-n, -1)| \leq |\mathcal{C}(-n, -1)| = -\mathcal{C}(-n, -1)$ if $n < J$,

$$(6) \quad \chi = 2 \quad \text{implies} \quad \lim_{n \rightarrow \infty} \mathbb{E}[|\mathcal{D}(-n, -1)|1_{J > n}] = 0.$$

5. On the event $X(0) \neq \boxed{\text{F}}$ the value $\mathcal{D}(0)$ is determined by $X(0)$ independently of $\mathcal{D}(-J, -1)$. Hence, by (5), the expectation of $\mathcal{D}(0)\mathcal{D}(-J, -1)$ restricted to this event is zero. On the event $X(0) = \boxed{\text{F}}$ we have $J = \phi(0)$ and hence $\mathcal{D}(0)$ has sign opposite that of $\mathcal{D}(-J, -1)$. Since $X(0) = \boxed{\text{F}}$ with probability $p/2$ and [even after conditioning on $X(0) = \boxed{\text{F}}$] we have $\mathbb{E}[|\mathcal{D}(-J, -1)|] = \chi$, this implies

$$(7) \quad \mathbb{E}[\mathcal{D}(0)\mathcal{D}(-J, -1)] = -\frac{\chi P}{2}.$$

6. On the event that $J < n$ the expectation of $\mathcal{D}(0)\mathcal{D}(-n, -J - 1)$ is zero, since $\mathcal{D}(-n, -J - 1)$ is independent of $\mathcal{D}(0)$ on this event [because $X(-J, 0)$ contains no $\boxed{\text{F}}$ symbols, and one can swap the roles of hamburgers and cheeseburgers in the word $X(-J)X(-J + 1) \cdots X(0)$ independently of the values of X_j for $j < -J$]. Thus the expectations of $\mathcal{D}(0)\mathcal{D}(-n, -1)$ and $\mathcal{D}(0)\mathcal{D}(-J, -1)$ are the same on the event $J < n$ (and trivially also the same on the event $J = n$). Thus,

$$(8) \quad \begin{aligned} &\mathbb{E}[\mathcal{D}(0)\mathcal{D}(-n, -1)] \\ &= \mathbb{E}[\mathcal{D}(0)\mathcal{D}(-J, -1)1_{J \leq n}] + \mathbb{E}[\mathcal{D}(0)\mathcal{D}(-n, -1)1_{J > n}]. \end{aligned}$$

By (6), the latter term in (8) tends to zero as $n \rightarrow \infty$ when $\chi = 2$. Thus,

$$(9) \quad \chi = 2 \quad \text{implies} \quad \lim_{n \rightarrow \infty} \mathbb{E}[\mathcal{D}(0)\mathcal{D}(-n, -1)] = -\frac{\chi P}{2} = -p.$$

Using translation invariance of the law of Y_j and (9), we obtain that

$$(10) \quad \begin{aligned} &\chi = 2 \quad \text{implies} \\ &\text{Var}[\mathcal{D}_n] = \sum_{i=1}^n \mathbb{E}[\mathcal{D}(Y_i)^2] + 2 \sum_{i=2}^n \mathbb{E}[\mathcal{D}(i)\mathcal{D}(1, i - 1)] = n - 2\frac{\chi P}{2}n + o(n), \end{aligned}$$

which in particular implies that $1 - \chi p \geq 0$ so that

$$(11) \quad \chi = 2 \quad \text{implies both} \quad p \leq 1/2 \quad \text{and} \quad \text{Var}[\mathcal{D}_n] = (1 - 2p)n + o(n).$$

7. Lemma 3.2 below states that the conclusion of (6) holds even if $\chi \neq 2$. For now, we note that the above calculations show that the conclusions of (9) and hence (10) remain true when $\chi \neq 2$ contingent on this claim:

$$(12) \quad \lim_{n \rightarrow \infty} \mathbb{E}[\mathcal{D}(-n, -1) | 1_{J > n}] = 0 \quad \text{implies} \quad \text{Var}[\mathcal{D}_n] = n - \chi pn + o(n).$$

From this and (3), we see the following:

$$(13) \quad \begin{aligned} &\text{Var}[\mathcal{D}_n] = o(n) \quad \text{and} \quad \lim_{n \rightarrow \infty} \mathbb{E}[|\mathcal{D}(-n, -1)| | 1_{J > n}] = 0 \quad \text{together imply} \\ &1 - \chi p = 0 \quad \text{and hence} \quad \chi = 1/p \quad \text{and} \quad p \geq 1/2. \end{aligned}$$

The following lemma will be proved later.

LEMMA 3.2. *If $\chi \neq 2$, then $\text{Var}[\mathcal{D}_n] = o(n)$ and $\lim_{n \rightarrow \infty} \mathbb{E}[|\mathcal{D}(-n, -1)| \times 1_{J > n}] = 0$.*

Lemma 3.2 states, in other words, that for each $p \in [0, 1]$, either the left statement in (11) (that $\chi = 2$) is true or the left statements in (13) (that $\text{Var}[\mathcal{D}_n] = o(n)$ and $\lim_{n \rightarrow \infty} \mathbb{E}[|\mathcal{D}(-n, -1)| | 1_{J > n}] = 0$) are true. We now claim that Lemma 3.1 is a consequence of (6), (11), (13), and Lemma 3.2. One hand Lemma 3.2 and (13) imply that if $p < 1/2$ then $\chi = 2$ and (11) gives the top row of Lemma 3.1. On the other hand (11) implies that if $p > 1/2$ then $\chi \neq 2$, and Lemma 3.2 and (13) give the bottom row. If $p = 1/2$, then the top and bottom row are equivalent, and we must have $\chi = 2$ [giving the desired result by (11)] since $\chi \neq 2$ would lead to a contradiction using Lemma 3.1 and (13).

In preparation for proving Lemma 3.2, we will derive several additional consequences of the assumption that $\chi \neq 2$ in Section 3.3.

3.2. *Two simple finite expectation criteria.* This section makes two simple observations that will be useful in Section 3.3. The first is a special case of what is sometimes called Wald’s identity. (We include a short proof for convenience in case the reader has not seen this before.)

LEMMA 3.3. *Let Z_1, Z_2, \dots be i.i.d. random variables on some measure space and ψ a measurable function on that space for which $\mathbb{E}[\psi(Z_1)]$ is well defined and finite. Let T be a stopping time of the process Z_1, Z_2, \dots with the property that $\mathbb{E}[T]$ is finite. Then $\mathbb{E}[\sum_{j=1}^T \psi(Z_j)]$ is also well defined and finite.*

PROOF. It is enough to consider the case that ψ is nonnegative (since we can write a general ψ as a difference of nonnegative functions). Since T is a stopping time, we know that for each fixed j , the value of Z_j is independent of the event

that $T \geq j$. Thus

$$\begin{aligned}
 \mathbb{E} \left[\sum_{j=1}^T \psi(Z_j) \right] &= \sum_{j=1}^{\infty} \mathbb{E}[\psi(Z_j) 1_{T \geq j}] \\
 (14) \qquad \qquad \qquad &= \sum_{j=1}^{\infty} \mathbb{E}[\psi(Z_j)] \mathbb{E}[1_{T \geq j}] = \mathbb{E}[\psi(Z_1)] \mathbb{E}[T]. \quad \square
 \end{aligned}$$

LEMMA 3.4. *Let Z_1, Z_2, \dots be i.i.d. random variables on some measure space and let \mathcal{Z}_n be a nonnegative-integer-valued process adapted to the filtration of the Z_n (i.e., each \mathcal{Z}_n is a function of Z_1, Z_2, \dots, Z_n) that has the following properties:*

1. Bounded initial expectation: $\mathbb{E}[\mathcal{Z}_1] < \infty$.
2. Positive chance to hit zero when close to zero: *For each $k > 0$, there exists a positive p_k such that $\mathbb{P}[\mathcal{Z}_{n+1} = 0 | Z_1, \dots, Z_k] \geq p_k$ on $\{\mathcal{Z}_n = k\}$.*
3. Uniformly negative drift when far from zero: *There exist positive constants C and c such that $\mathbb{E}[\mathcal{Z}_{n+1} - \mathcal{Z}_n | Z_1, \dots, Z_k] \leq -c$ on $\{\mathcal{Z}_n \geq C\}$.*
4. Bounded expectation when near zero: *There further exists a constant b such that $\mathbb{E}[\mathcal{Z}_{n+1} | Z_1, \dots, Z_n] < b$ on $\{\mathcal{Z}_n < C\}$.*

Then

$$\mathbb{E}[\min\{n : \mathcal{Z}_n = 0\}] < \infty.$$

PROOF. The uniformly negative drift assumption implies that the quantity $Y_n = \mathcal{Z}_n + cn$ is a supermartingale until time $K = \min\{n \geq 0 : \mathcal{Z}_n < C\}$. In particular, this fact together with the optional stopping theorem implies that as a function of k the quantity $\mathbb{E}[Y_{K \wedge k}]$ is bounded above independently of k . On the event $K \geq k$ we have $Y_k \geq ck$. Since $Y_{K \wedge k}$ is non-negative, this implies that $ck\mathbb{P}[K \geq k] \leq \mathbb{E}[Y_{K \wedge k}]$. In particular, this implies that $\lim_{k \rightarrow \infty} \mathbb{P}[K \geq k] = 0$ so that K is a.s. finite.

Next, by the optional stopping theorem for non-negative supermartingales, $\mathbb{E}[Y_K] \leq \mathbb{E}[Y_1]$ which implies $\mathbb{E}[cK] \leq \mathbb{E}[Y_1]$ and $\mathbb{E}[K] \leq \mathbb{E}[Y_1]/c$. Now let T_1, T_2, \dots denote the successive times when $\mathcal{Z}_n < C$. Then we claim that a similar analysis implies that $\mathbb{E}[T_{k+1} - T_k | Z_1, Z_2, \dots, Z_{T_k}] \leq \beta$ for some positive constant β . To carry out this analysis, we will write $\tilde{Y}_n = \mathcal{Z}_{T_k+n} + cn$, and note that the same argument as above (using probability conditioned on Z_1, Z_2, \dots, Z_{T_k} and noting that $\mathbb{E}[\tilde{Y}_1 | Z_1, Z_2, \dots, Z_{T_k}] \leq b + c$) shows that

$$\mathbb{E}[T_{k+1} - T_k | Z_1, Z_2, \dots, Z_{T_k}] \leq \mathbb{E}[\tilde{Y}_1 | Z_1, Z_2, \dots, Z_{T_k}]/c \leq (C + c)/c =: \beta.$$

Next, recall that whenever $\mathcal{Z}_n < C$, there is a probability of at least $\delta = \min\{p_k : k \leq C\} > 0$ that $\mathcal{Z}_{n+1} = 0$. This implies that if $S := \min\{n : \mathcal{Z}_n = 0\}$ then $\mathbb{P}[S >$

$T_k] \leq \delta^{k-1}$. Since

$$(15) \quad \mathbb{E}[S] \leq E[T_1] + E \left[\sum_{k=1}^{\infty} (T_{k+1} - T_k) 1_{S > T_k} \right]$$

$$(16) \quad \leq \mathbb{E}[K] + \sum_{k=1}^{\infty} \mathbb{P}[T_k > S] \mathbb{E}[T_{k+1} - T_k | S > T_k]$$

$$(17) \quad \leq \mathbb{E}[K] + \sum_{k=1}^{\infty} \delta^{k-1} \beta < \infty. \quad \square$$

3.3. *Excursion words.* In Section 3.1, we considered the random word $X(-J, -1)$, where J was the smallest integer for which this word had at least one \textcircled{C} or \textcircled{H} symbol. We found that the expected word length $\mathbb{E}[|X(-J, -1)|]$ was a constant $\chi \in [1, 2]$. In this section, we will consider different words and show that they all have finite expected length provided that $\chi < 2$.

Suppose that K is the smallest $k \geq 0$ for which $C_{k+1} = C(1, k + 1) < 0$. If $X(1) \in \{\textcircled{C}, \textcircled{H}, \textcircled{F}\}$ then $K = 0$. Otherwise, K is a positive value for which $C_K = C(1, K) = 0$, so that $X(1, K)$ is “balanced” in the sense that it has the same number of burgers as orders. Call $E = X(1, K)$ the *excursion word* beginning at time zero (writing $E = \emptyset$ if $K = 0$). We make a few observations about E :

1. E a.s. contains no \textcircled{F} symbols. [Indeed, one can check inductively that $X(1, k)$ contains no \textcircled{F} symbols for any $1 \leq k \leq K$.]
2. If $p = 1$, then $E = \emptyset$ almost surely.
3. The law of K is independent of $p \in [0, 1]$. Its law is that of the number of steps taken by a simple random walk on \mathbb{Z} , started at 0, before it first hits -1 . (In particular, $\mathbb{E}[K] = \infty$.)

We denote by V_i the symbol corresponding to the i th record minimum of C_n , counting forward from zero, if i is positive, and the $-i$ th record minimum of C_n , counting backward from zero, if i is negative (see Figure 4). We will introduce notation to describe the locations of the V_i in the proof of Lemma 3.5.

Denote by E_i the *reduced form* of the word strictly in between the symbols corresponding to V_i and V_{i-1} (if i is positive) or between V_i and V_{i+1} (if i is negative), where V_0 is for the moment formally taken to be the “zero point” located between the symbols $X(0)$ and $X(1)$ (see Figure 4). We do not define E_i for $i = 0$. Note that the V_i are i.i.d. equal to \textcircled{H} or \textcircled{C} (each with probability 0.5) when $i < 0$ and equal to \textcircled{H} , \textcircled{C} , or \textcircled{F} [with probabilities $(1 - p)/2$, $(1 - p)/2$ and p] when $i > 0$. The E_i are i.i.d. excursion words, each with same law as the E described above, and are independent of the V_i . (The fact that E_1 and E_{-1} are identically distributed follows from the fact that an excursion of the simple random walk C_n is equivalent in law to its time reversal. Indeed, once we condition on the trajectory of C_n over




FIG. 4. Possible graph of C_n as a function of n . Up-going edges correspond to burgers, down-going edges to orders. The V_i correspond to edges that reach record minima when one starts from zero and moves in one (positive or negative) direction. The E_i correspond to the excursions in between the V_i . [Recall a slight asymmetry in the notation: the first edge to the right of vertex labeled 0 corresponds to $X(1)$ while the first edge to the left corresponds to $X(0)$.]

an excursion, we can sample the corresponding symbols by tossing independent coins to replace each upward step with an \textcircled{C} or \textcircled{H} and each downward step with a \boxed{C} , \boxed{H} , or \boxed{F} .)

LEMMA 3.5. Suppose that p is such that the χ defined in Section 3.1 satisfies $\chi < 2$. Then the expected word length $\mathbb{E}[|E|]$ is finite. Moreover, the expected number of symbols in E of each of the four types in $\{\textcircled{C}, \textcircled{H}, \boxed{C}, \boxed{H}\}$ is $\mathbb{E}[|E|]/4$.

PROOF. Since E has the same number of symbols in $\{\textcircled{C}, \textcircled{H}\}$ as in $\{\boxed{C}, \boxed{H}\}$ (recall that E has no \boxed{F} symbols), the second statement is immediate from the first by symmetry. To prove that $\mathbb{E}[|E|] < \infty$, it suffices to show that the expected number of burgers in the word E_{-1} is finite.

To this end, let J_1 be the smallest nonnegative integer j for which $X(-j, 0)$ has a \textcircled{C} or \textcircled{H} symbol. This is the same as the J defined in Section 3.1 except that we count from 0 instead of from -1 . [The law of the word $X(-J_1, 0)$ is the same as the law of the word $X(-J, -1)$ in Section 3.1. We count from 0 here because $X(0)$ is the rightmost symbol involved in the definition of E_{-1} and V_{-1} , see Figure 4.] Then let J_2 be the smallest value greater than J_1 for which $X(-J_2, -J_1 - 1)$ has a \textcircled{C} or \textcircled{H} symbol, and so forth. The words $X(-J_{k+1}, -J_k - 1)$ are i.i.d. and all have the same laws as $X(-J, -1)$ (and each a.s. has length at least one). In particular, if $\chi < 2$, then we have by (2) that $\mathbb{E}[\mathcal{C}(-J_{k+1}, -J_k - 1)] = 2 - \chi > 0$ for each k .

This implies that (if we write $J_0 = -1$) the sum

$$(18) \quad \sum_{k=1}^m \mathcal{C}(-J_k, -J_{k-1} - 1) - (2 - \chi)m$$

is a martingale indexed by m . Note that the sum $\sum_{k=1}^m \mathcal{C}(-J_k, -J_{k-1} - 1)$ can increase by at most 1 at each increment. Let M be the smallest m for which

$\sum_{k=1}^m \mathcal{C}(-J_k, -J_{k-1} - 1) = 1$. The fact that $\sum_{k=1}^m \mathcal{C}(-J_k, -J_{k-1} - 1) = 1$ is a sum of i.i.d. integer-valued random variables, each with expectation greater than zero, together with the law of large numbers implies that the sum tends to ∞ as $m \rightarrow \infty$. Since each of the random variables in the sum is at most 1 we conclude that $M < \infty$ a.s. Since the martingale (18) is bounded above up to time M , the optional stopping theorem for martingales bounded above implies that the expected value of (18) at this stopping time is at least 0, which implies that $\mathbb{E}[1 - (2 - \chi)M] \geq 0$ and hence $\mathbb{E}[M] \leq 1/(2 - \chi)$. Since $\mathcal{C}(-J_M, 0) = 1$, it follows that $-J_M$ is left of the location of V_{-1} (since the latter corresponds to the *first* time the count reaches 1, counting left from zero). Since the number of burgers in $X(-j, 0)$ is increasing in j , it follows that $X(-J_M, 0)$ has at least as many burgers as $\overline{V_{-1}E_{-1}}$. The expectation of the number of burgers in $X(-J_M, 0)$ is finite [since $X(-J_M, 0)$ has at most M burgers], and hence so is the expectation of the latter. \square

The lemmas that follow will involve several sequences related to the J_m defined above (each defined for $m \geq 1$). For convenience, we define them here:

1. *mth empty order stack*: O_m is the m th smallest value of $j \geq 0$ with the property that $X(-j, 0)$ has an empty order stack. We have not proved that O_1 is a.s. finite, but if it is then we may observe that the words $X(-O_m, -O_{m-1} - 1)$ are i.i.d. (if we formally set $O_0 = -1$) and that each has an empty order stack.
2. *mth empty burger stack*: B_m is the m th smallest value of $j \geq 1$ with the property that $X(1, j)$ has an empty burger stack. We have not proved that B_1 is a.s. finite, but if it is, then we may observe that the words $X(B_{m-1} + 1, B_m)$ are i.i.d. (if we formally set $B_0 = 0$) and that each has an empty burger stack.
3. *mth left record minimum*: L_m is the smallest value of $j \geq 0$ for which $\mathcal{C}(-j, 0) = m$. Thus $X(-L_m, 0) = \overline{V_{-m}E_{-m} \cdots V_{-1}E_{-1}}$ and $X(-L_m, -L_{m-1} - 1) = \overline{V_{-m}E_{-m}}$ (if we formally set $L_0 = -1$).
4. *mth right record minimum*: R_m is the smallest value of $j \geq 1$ for which $\mathcal{C}(1, j) = -m$. Thus $X(1, R_m) = \overline{E_1V_1 \cdots E_mV_m}$ and $X(R_{m-1} + 1, R_m) = \overline{E_mV_m}$ (if we formally set $R_0 = -1$).
5. *mth hamburger-order-free left minimum*: L_m^H is the m th smallest value of $j \geq 0$ with the property that $X(-L_j^H, 0)$ has no hamburger orders and $L_j^H = L_{j'}$ for some j' [i.e., $-L_j^H$ corresponds to a new record minimum of $\mathcal{C}(\cdot, 0)$]. As in the cases above, the words $X(-L_j^H, -L_{j-1}^H - 1)$ are i.i.d., and each one is a product of some sequence of words of the form $X(-L_{j'}, -L_{j'-1} - 1)$.
6. *mth hamburger-free right minimum*: similarly, R_m^H is the m th smallest value of $j \geq 1$ with the property that $X(1, R_m^H)$ has no hamburgers and $R_j^H = R_{j'}$ for some j' .

LEMMA 3.6. *The following are equivalent:*

1. $\mathbb{E}[|E|] < \infty$.

2. $\mathbb{E}[|X(-L_1^H, 0)|] < \infty$.
3. $\mathbb{E}[|X(1, R_1^H)|] < \infty$.
4. O_1 is a.s. finite and $\mathbb{E}[|X(-O_1, 0)|] < \infty$.
5. B_1 is a.s. finite and $\mathbb{E}[|X(1, B_1)|] < \infty$.

PROOF. 1 implies 2: Assume 1. Recall that the words E_{-j} are i.i.d., that each a.s. has no $\boxed{\text{F}}$ symbols and as many orders as burgers (possible example: $\boxed{\text{C}}\boxed{\text{H}}\boxed{\text{H}}\boxed{\text{C}}\boxed{\text{C}}\boxed{\text{C}}\boxed{\text{C}}\boxed{\text{H}}$), that the expected number of symbols of each of the types $\{\boxed{\text{C}}, \boxed{\text{H}}, \textcircled{\text{C}}, \textcircled{\text{H}}\}$ in E_{-1} is finite, and that these expectations are equal. Recall also that the V_j (for $j < 0$) are independent i.i.d. samples from $\{\textcircled{\text{C}}, \textcircled{\text{H}}\}$. Let $H(m)$ be the number of hamburger orders in $X(-L_m, 0) = \overline{V_{-m}E_{-m} \cdots V_{-1}E_{-1}}$. Note that for any $m > 1$ we have that

$$(19) \quad H(m) = \max\{H(m - 1) - h_m, 0\} + o_m,$$

where h_m is the number of hamburgers in $\overline{V_{-m}E_{-m}}$ and o_m is the number of hamburger orders in $\overline{V_{-m}E_{-m}}$. Note that the pair (h_m, o_m) is independent of $H(m - 1)$. Note also that $\mathbb{E}[h_m] = \mathbb{E}[o_m] + 0.5$ (since the expected number of hamburger orders in E equals the expected number of hamburgers and $V_{-m} = \textcircled{\text{H}}$ with probability 0.5).

Now we can rewrite (19), using the fact that $\max\{A, 0\} = A - 1_{A < 0}A$, as

$$H(m) - H(m - 1) = o_m - h_m + (h_m - H(m - 1))1_{\{H(m-1) - h_m < 0\}},$$

and $\mathbb{E}[(h_m - k)1_{h_m > k}] \leq \mathbb{E}[h_m 1_{h_m > k}] \downarrow 0$ as $k \rightarrow \infty$ by assumption. Thus, there exists a $C > 0$ such that for all $k > C$ we have $\mathbb{E}[H(m) - H(m - 1) | H(m - 1) = k] \leq -0.4$. It follows from Lemma 3.4, applied with $\mathcal{Z}_m = H(m)$ and $Z_m = \overline{V_{-m}E_{-m}}$, that the expected number of $V_{-j}E_{-j}$ sequences concatenated to produce $X(-L_1^H, 0)$ is finite, and it then follows from 1 and Lemma 3.3 that the expected sum of the lengths of these words is finite.

1 implies 3: The argument used to show that 1 implies 2 applies almost verbatim here, using R_m in place of L_m .

2 implies 4: We use the argument used to show 1 implies 2. In that case, we concatenated i.i.d. words of the form $X(-L_m, -L_{m-1} - 1) = \overline{V_{-m}E_{-m}}$ until we produced a word with no hamburger orders, which was equal to $X(-L_1^H, 0)$. In this case, we start by concatenating the i.i.d. words of the form $X(-L_m^H, -L_{m-1}^H - 1)$, each of which has no orders of type $\boxed{\text{H}}$ or $\boxed{\text{F}}$ (possible example: $\boxed{\text{C}}\boxed{\text{C}}\boxed{\text{C}}\boxed{\text{C}}\textcircled{\text{C}}\textcircled{\text{H}}\textcircled{\text{C}}\textcircled{\text{C}}\textcircled{\text{C}}\textcircled{\text{H}}$), and we continue until we produce a word with no orders at all. This word is necessarily at least as long as $X(-O_1, 0)$ [since the burger stack is increasing in time, and $X(-O_1, 0)$ corresponds to the first time the order stack is empty], so it suffices to bound its expected length.

We are assuming that the expected length of $X(-L_1^H, 0)$ is finite; since this word has at least as many burgers as $X(-L_1, 0)$, this implies that the expected

number of burgers in $X(-L_1, 0)$ is also finite, and since the number of orders in $X(-L_1, 0)$ is one less than the number of burgers, we find that the expected total length of $X(-L_1, 0)$ is finite, as is the expected total length of E_{-1} . That is, we have established that 2 implies 1. We may therefore use the fact (established at the end of the 1 implies 2 argument) that the expected number of $V_{-j}E_{-j}$ sequences concatenated to produce $X(-L_1^H, 0)$ is some finite number k .

Note that the expected net number of cheeseburgers in $X(-L_m, -L_{m-1} - 1) = V_{-m}E_{-m}$ (i.e., the expected number of \textcircled{C} symbols minus the number of \boxed{C} symbols) is $1/2$, since this net number is additive when we concatenate words with no \boxed{F} symbols. We similarly conclude by (14) that the expected net number of cheeseburgers in $X(-L_1^H, 0)$ is $k/2$, where k is as defined just above.

As we concatenate i.i.d. copies of the word $X(-L_1^H, 0)$, the number of \boxed{C} symbols in the reduced form of the concatenated word is a Markov chain on \mathbb{Z}_+ . The above observations imply that the expected change in this Markov chain value, during a step that starts at position j , is a value that tends to $-k/2$ as $j \rightarrow \infty$. For any j , this expected change is bounded above by the finite expected length of $X(-L_1^H, 0)$. Thus, Lemma 3.4, implies that the expected number of steps until this chain reaches zero is finite. Similarly, Lemma 3.3 then implies that the expected total length of the words concatenated is finite.

3 implies 5: This is essentially the same as the proof that 2 implies 4.

Either 4 or 5 implies 1: Note that the number of burgers in $X(-O_1, 0)$ is at least the number in E_{-1} , and the number of orders in $X(1, B_1)$ is at least the number in E_1 . Thus if either $\mathbb{E}[|X(-O_1, 0)|] < \infty$ or $\mathbb{E}[|X(1, B_1)|] < \infty$, then (recalling that the expected number of orders in E equals the expected number of burgers in E) we have $\mathbb{E}[|E|] < \infty$. \square

LEMMA 3.7. *If $\mathbb{E}[|E|] < \infty$, then the limit, as $n \rightarrow \infty$, of the fractions of \textcircled{H} symbols among the top n elements in $X(-\infty, 0)$ is equal to 0.5 almost surely. Similarly, as n tends to infinity, the fraction of \boxed{C} , \boxed{H} and \boxed{F} symbols among the leftmost n elements of $X(1, \infty)$ tend almost surely to constants (the first two equal, by symmetry). These constants are all positive (except in case $p = 1$, where one has only \boxed{F} symbols and no \boxed{C} or \boxed{H} symbols).*

On the other hand, if $\mathbb{E}[|E|] = \infty$ then the limit as $n \rightarrow \infty$ of the fraction of \boxed{F} symbols among the leftmost n elements of $X(1, \infty)$ tends almost surely to zero.

PROOF. Consider the sequence of words $X(-k, 0)$, indexed by positive k . As we have noted before, the words $X(-O_m, -O_{m-1} - 1)$ are i.i.d., each consisting entirely of burgers, and $X(-\infty, 0)$ is the concatenation of these all-burger words. Lemma 3.6 implies that if $\mathbb{E}[|E|] < \infty$, then the lengths of these burger sequences have finite expectation, and the law of large numbers then implies the number of hamburgers and cheeseburgers in $X(-O_m, 0)$ are both [up to $o(m)$ errors] given by constant multiples of m , almost surely. This implies the first statement in the

lemma. The proof of the second statement is analogous, using B_m instead of O_m . To see that the constants are positive, it suffices to observe that $X(1, B_1)$ has some positive probability of containing an order of each of the three order types.

For the final statement, we note that an \boxed{F} symbol can be added to the order stack of the sequence $X(1, k)$ only at times when the burger stack is empty. The number of \boxed{F} symbols in $X(1, B_m)$ can grow as a function of m , but it can grow by at most 1 each time that m increases by 1. If $\mathbb{E}[|E|] = \infty$, then Lemma 3.6 and the law of large numbers imply that the number of orders in $X(1, B_m)$ a.s. grows faster than any constant times m , while the number of \boxed{F} symbols grows like a constant times m . \square

The following proposition is not needed for the proof of Theorem 2.5, but we include it because it will be interesting from the point of view of random planar maps. (In a sense, it will imply that the infinite random surface models one obtains when $p > 1/2$ a.s. have infinitely many small bottlenecks surrounding any given point.)

PROPOSITION 3.8. *Let K be the infimum over the set of positive numbers such that there exist $m_- < 0$ and $m_+ > 0$ such that:*

1. $X(m_-, 0)$ is a word with no orders and K burgers.
2. $X(1, m_+)$ is a word with no burgers and K orders.
3. $X(m_-, m_+) = \emptyset$.

If $\mathbb{E}[|E|] < \infty$, then K is a.s. finite and $\mathbb{E}[K] < \infty$.

PROOF. Assume $\mathbb{E}[|E|] < \infty$. As noted in the proof of Lemma 3.7, the gaps between $a_m := |X(-O_m, 0)|$ (as m increases) and between $b_m := |X(1, B_m)|$ (as m increases) are i.i.d. positive random variables with finite expectation (and clearly every possible positive gap length has positive probability).

Write $\alpha_k = k - \max\{a_m : a_m \leq k\}$ and $\beta_k = k - \max\{b_m : b_m \leq k\}$. Both α_k and β_k can be understood as ergodic Markov chains. To see this, note that if $\alpha_k = j$, then $\alpha_{k+1} \in \{j + 1, 0\}$ and $\mathbb{P}[\alpha_{k+1} = 0 | \alpha_k = j] = \mathbb{P}[a_1 = j + 1 | a_1 > j]$. The transition kernel for β_k can be written similarly. Since a . and b . each have increments with finite expectation, the corresponding ergodic Markov chains assume the value 0 (at stationarity) with probability one over the corresponding expectation. Since α . and β . are independent, the pair (α, β) is also a Markov chain that (at stationarity) assumes the value $(0, 0)$ with positive probability.

Call a nonnegative integer M a “match up length” if there exists an i such that $a_i = |X(-O_i, 0)| = M$ and a j such that $b_j = |X(1, B_j)| = M$. The match up lengths are a random subset of \mathbb{Z}_+ and correspond to the k for which $(\alpha_k, \beta_k) = (0, 0)$. The discussion above implies that the gaps between successive match up lengths are i.i.d. random variables with finite expectation. Denote by M_k the k th such match up length.

For each M_k , we have a word $X(-O_i, 0)$ of length M_k comprised entirely of burgers and a word $X(1, B_j)$ of length M_k , comprised entirely of orders. We claim that $X(-O_i, B_j)$ is a word that consists of a sequence of orders of one type (either \boxed{C} or \boxed{H}) followed by a same-length sequence of burgers of opposite type (either \textcircled{C} or \textcircled{H}), for example,

$$\boxed{C}\boxed{C}\boxed{C}\boxed{C}\boxed{C}\textcircled{H}\textcircled{H}\textcircled{H}\textcircled{H}\textcircled{H}.$$

[This can be seen by starting with the all-burger word $X(-O_i, 0)$ and multiplying on the right by the symbols in $X(1, B_j)$ one at a time. The only way there can be orders in $X(-O_i, B_j)$ is if all of the burgers of one type are consumed before this process terminates, and in that case only orders corresponding to that burger type can appear in $X(-O_i, B_j)$.] When we shift from k to $k + 1$, we multiply this word on the left and right by random words W_1 and W_2 . We know that $\mathcal{C}(W_1) + \mathcal{C}(W_2) = 0$. By symmetry, the expected number of symbols of type \textcircled{H} in this pair of words equals the expected number of symbols of type \textcircled{C} . (These expectations are finite by Lemma 3.6.) Similarly for \boxed{H} and \boxed{C} symbols. The expected number of \boxed{F} symbols is some positive constant a .

If the length of $X(-O_i, B_j)$ is long enough, then the probability that either W_1 or W_2 is longer than half of that length is close to zero. Thus, conditioned on the length being at least some constant value, the expected change in length as one goes from k to $k + 1$ is close to $-a/2$. It then follows from Lemma 3.4 that the expected value of the smallest k for which this length is zero is finite, and the result then follows from Lemma 3.3. \square

3.4. *Some monotonicity observations.* In this section, it will be convenient to fix a semi-infinite stack present at time zero [the stack $X(-\infty, 0)$] and also to relax the assumption in Definition 2.4 that $(\mathcal{D}_n, \mathcal{C}_n) = (0, 0)$. To give some notation for this, let S_0 be a semi-infinite stack of burgers containing infinitely many burgers of each type *together* with a corresponding diagram in \mathbb{R}^2 such as the one given in Figure 2. The diagram is determined by the burger stack once we know the location of the uppermost vertex (the *tip*) in Figure 2, which we now allow to be *any* lattice point of \mathbb{Z}^2 whose coordinate sum is even. Slightly abusing notation, we will write $(\mathcal{D}_0, \mathcal{C}_0)$ for the location of the tip of S_0 (which, in this section only, we do not require to be at the origin).

Given S_0 , we can generate a sequence of stacks S_1, S_2, \dots in the usual manner by applying the moves in Figure 2 that correspond to $X(1), X(2), \dots$. For $n > 0$, we write $(\mathcal{D}_n, \mathcal{C}_n) := (\mathcal{D}_0 + \mathcal{D}(1, n), \mathcal{C}_0 + \mathcal{C}(1, n))$ as before, so that $(\mathcal{D}_n, \mathcal{C}_n)$ is the location of the tip of S_n .

Given any such embedded stacks S and \tilde{S} , we write $S \leq \tilde{S}$ if the tips of the path lie on the same horizontal line and the path describing S (as in Figure 2) lies to the left of the path describing \tilde{S} —that is, every horizontal line intersects the S path at a point equal to or left of where it intersects the \tilde{S} path, see Figure 5. The following lemma is immediate from an inspection of the cases in Figure 2.




FIG. 5. Upper portion of embedded stacks $S \leq \tilde{S}$ (partially overlapping).

LEMMA 3.9. Suppose embedded stacks $S_0 \leq \tilde{S}_0$ are fixed and that S_1, S_2, \dots and $\tilde{S}_1, \tilde{S}_2, \dots$ are generated from S_0 and \tilde{S}_0 respectively using the same sequence $X(1), X(2), \dots$ as in Figure 2. Then $S_n \leq \tilde{S}_n$ for all $n > 0$.

We can use this result to deduce the following lemma.

LEMMA 3.10. Fix $N > 0$ and S_0 and let $X(1), X(2), \dots$ be chosen from μ . Then $\mathbb{E}[\mathcal{D}_N | X(j) : j \leq k]$ is a martingale in k with increments of magnitude at most 2 (which obtains the value \mathcal{D}_N at $k = N$). A similar result holds when one further conditions on the value of the sequence C_n for $n \in \{1, 2, \dots, N\}$: namely, $\mathbb{E}[\mathcal{D}_N | X(j) : j \leq k, C_j : 0 \leq j \leq N]$ is a martingale in k with increments of magnitude at most 2.

The lemma implies, in particular, that if we fix any choice of S_0 and fix $N > 0$ and then sample the sequence $X(1), X(2), \dots$ chosen from μ , the variance of \mathcal{D}_N is at most $4N$. (The conditional variance of C_N is exactly N .)

PROOF OF LEMMA 3.10. We claim that changing the value $X(j)$ for a single $j \in \{1, 2, \dots, N\}$ (while leaving the other values fixed) always changes \mathcal{D}_N by at most 2. Both statements in the lemma are immediate consequences of this. We will establish the claim using Lemma 3.9.

First, if we change $X(j)$ in a way that does not affect C_j [i.e., we change $X(j)$ from one burger type to another, or from one order type to another] then it is clear from Figure 2 that the modified S_j obtained lies between the original S_j and the original translated by 2 units (either left or right). It follows from Lemma 3.9 that the change to $X(j)$ alters the tip of the final stack S_N by at most 2 units. If we change $X(j)$ in a way that *does* affect C_j , then a variation of this argument still applies. Suppose the modified S_j has a tip two units higher vertically than the original. Then if we translate the modified S_j *down* by two units, it will lie within two units to the left or right of the original S_j (again, by inspection of Figure 2), and one can then apply Lemma 3.9 as before to show that this remains true after both stacks evolve for $N - j$ steps using the same $X(j + 1) \cdots X(n)$ sequence. □

The following is a fairly simple and standard observation about the tail behavior of martingales with small increments. It follows for example, from Lemma 2 of [2]. (We include a short proof here for completeness, but this can be skipped.)

LEMMA 3.11. *Let β_j be a martingale in j whose increments have magnitude at most 1, with $\beta_0 = 0$. Then for each real $a > 0$ and integer $n > 0$ we have*

$$P\left(\max_{j \in \{1, 2, \dots, n\}} n^{-1/2} \beta_j \geq a\right) \leq e^{-a+1/2}.$$

PROOF. First, we observe that $\frac{e^b + e^{-b}}{2} \leq e^{b^2/2}$ for all b . [Indeed, by Taylor expansion, the left side is $1 + b^2/2! + b^4/4! + b^6/6! + \dots$ while the right side is $1 + b^2/(2^1 1!) + b^4/(2^2 2!) + b^6/(2^3 3!) + \dots$.] Thus, if the increments of β_j were exactly ± 1 (each with probability $1/2$), with $\beta_0 = 1$, then we would have $\mathbb{E}[e^{b\beta_1}] \leq e^{b^2/2}$ and $\mathbb{E}[e^{b\beta_1 - b^2/2}] \leq 1$. We claim that the same is true if we allow for $\beta_1 \notin \{-1, 1\}$ and insist only that $|\beta_1| \leq 1$ a.s. and $\mathbb{E}\beta_1 = 0$. (Indeed, one may first choose an increment β_1 , and then, given this, choose a “modified increment” $\beta'_1 \in \{-1, 1\}$ using a biased coin whose probability is chosen so that conditional expectation of the modified increment is β_1 . Jensen’s inequality implies $\mathbb{E}[e^{b\beta_1}] \leq \mathbb{E}[e^{b\beta'_1}]$.) The argument above shows more generally that $e^{b\beta_j - jb^2/2}$ is a supermartingale indexed by j .

Taking $b = n^{-1/2}$ produces a supermartingale in j whose value at time $j = 0$ is 1 and whose value at general j time is $e^{n^{-1/2}\beta_j - j/(2n)}$. The expectation of the latter quantity is at most 1, by the optional stopping theorem, so the probability that it ever reaches $e^{a-1/2}$ is at most $e^{-a+1/2}$. To conclude, we note that $n^{-1/2}\beta_j \geq a$ implies $e^{n^{-1/2}\beta_j - j/(2n)} \geq e^{a-1/2}$ for $j \in \{1, 2, \dots, n\}$. \square

LEMMA 3.12. *Fix any $p \in [0, 1]$. Fix the time zero stack S_0 . Then there are positive constants C_1 and C_2 such that, for any choice of S_0 , $a > 0$, and $n \geq 1$, the probability that $|\mathcal{D}_j| > a\sqrt{n}$ for some $j \in \{0, 1, 2, \dots, n\}$ is at most $C_1 e^{-C_2 a}$.*

PROOF. Without loss of generality, assume $C_0 = \mathcal{D}_0 = 0$. To sample $X(1), X(2), \dots$ we may first sample the sequence $\mathcal{C}_1, \mathcal{C}_2, \dots$ and then conditioned on that choose the types of the burgers and orders (which in turn determine the \mathcal{D}_j sequence). Write $M = a\sqrt{n}/8$ and note that

$$(20) \quad \mathbb{P}\left\{\max_{j \in \{1, 2, \dots, n\}} |\mathcal{C}_j| > M\right\} \leq 2e^{-a/8+1/2}$$

by Lemma 3.11. Thus, we may restrict attention to the event \mathcal{E}_n that this does not occur, in which case we are conditioning on values of \mathcal{C}_j that satisfy $|\mathcal{C}_j| \leq M$ for $j \in \{1, 2, \dots, n\}$.

Using Lemma 3.10, we have that

$$(21) \quad m_j := \mathbb{E}[\mathcal{D}_n | X(i) : i \leq j, \mathcal{C}_i : 0 \leq i \leq n]$$

is a martingale in j for $j \in \{0, 1, \dots, n\}$ with increments of size at most 2. The value of \mathcal{D}_j depends on the stack at time j and does not depend on any of the

edges in S_0 that lie below $-M$. This is because if one considers two *different* choices of S_0 which agree above level $-M$, and then takes the *same* choice of symbols $X(0), X(1), X(2), \dots$ whose values imply $|C_j| \leq M$ for $j \in \{1, 2, \dots, n\}$, then by induction the updated stack S_j will look the same above level $-M$ for each j between 0 and n , and since all \boxed{F} symbols get matched to burgers above level $-M$, we may conclude that (on the event \mathcal{E}_n) \mathcal{D}_j does not depend on the choice of S_0 below level $-M$. We can also rewrite (21) as $m_j = \mathbb{E}[\mathcal{D}_n | S_j, C_i : 0 \leq i \leq n]$.

Now fix some value $j \in \{1, 2, \dots, n\}$ and let \tilde{S}_j be the stack obtained by swapping the burgers in S_j above level $-M$ and subsequently translating the stack by $4M$ units to the left (i.e., translating the stack by the amount that ensures that the corresponding \mathcal{D}_j value is translated by $4M$, but C_j remains the same). Then we see that \tilde{S}_j lies strictly to the left of S_j . One can then see inductively (by Lemma 3.9) that this will continue to be the case if we evolve both stacks in parallel by adding the same symbols X_{j+1}, X_{j+2}, \dots to each. Let $\tilde{\mathcal{D}}_j$ be the discrepancy of \tilde{S}_j , so that (by definition) we have $\tilde{\mathcal{D}}_j = \mathcal{D}_j - 4M$. Let $\tilde{\mathcal{D}}_i$ for $j \leq i \leq n$ be obtained by letting the stack evolve by the addition of the symbols $X_j + 1, X_j + 2, \dots, X_n$. Set $\tilde{m}_j = \mathbb{E}[\tilde{\mathcal{D}}_n | \tilde{S}_j; C_i, 0 \leq i \leq n]$, which is also $\mathbb{E}[\tilde{\mathcal{D}}_n | S_j, C_i : 0 \leq i \leq n]$. Now observe by symmetry that on the event \mathcal{E}_n we have

$$(22) \quad \tilde{\mathcal{D}}_j - \tilde{m}_j = -(\mathcal{D}_j - m_j).$$

By the fact (noted above) that \tilde{S}_i is left of S_i for $j \leq i \leq n$ one has $\tilde{\mathcal{D}}_n \leq \mathcal{D}_n$ and therefore $\tilde{m}_j \leq m_j$, and we also recall that $\tilde{\mathcal{D}}_j = \mathcal{D}_j - 4M$. Rewriting the LHS of (22) using these two facts, we get

$$\mathcal{D}_j - 4M - m_j \leq -\mathcal{D}_j + m_j,$$

which implies $\mathcal{D}_j - m_j \leq 2M$ on \mathcal{E}_n .

An analogous statement holds if we define \tilde{S}_j using translation to the right instead of to the left. Thus, on the event \mathcal{E}_n , we have $|\mathcal{D}_j - m_j| \leq 2M = \frac{a}{4}\sqrt{n}$ for all $j \in \{0, 1, \dots, n\}$. Thus, in order to have $|\mathcal{D}_j - \mathcal{D}_0| \geq a\sqrt{n}$ for some $j \in \{1, 2, \dots, n\}$ we must have $|m_j - m_0| \geq a\sqrt{n}/2$. Applying Lemma 3.11 to $m_j/2$, we find that the conditional probability (given \mathcal{E}_n) that $|m_j - m_0|$ exceeds $\sqrt{na}/2$ on the interval $\{1, 2, \dots, n\}$ is at most $e^{-a/4+1/2}$. This combined with (20) gives the lemma. \square

LEMMA 3.13. *The length of $X(1, n)$ is typically of order \sqrt{n} or smaller regardless of p . More precisely, there are positive constants C_1 and C_2 (independent of n and $a > 0$) such that*

$$P\left(\sup_{i \in \{1, 2, \dots, n\}} |X(1, i)| > a\sqrt{n}\right) \leq C_1 e^{-C_2 n}.$$

PROOF. Fix an initial stack S_0 to be alternating between \textcircled{C} and \textcircled{H} . Then observe that if \mathcal{D}_k and \mathcal{C}_k both fluctuate by at most $a\sqrt{n}/5$, as k ranges from 1 to n , then no burger on the stack S_0 strictly below height $-2a\sqrt{n}/5$ will have been consumed during the first n steps. (If the first burger below that height to be consumed were consumed at step j , then all of the burgers above it in the stack at step j —let k be the number of such burgers—would have to be of the same type. This would require that either $k < a\sqrt{n}/5$, in which case $|\mathcal{C}_j|$ would have to exceed the assumed bound, or that $k \geq a\sqrt{n}/5$, in which case $|\mathcal{D}_j|$ would have to exceed the assumed bound, since all the burgers below the top k must strictly alternate between \textcircled{C} and \textcircled{H} .) Thus, on this event the total number of orders in $X(1, i)$ is less than $2a\sqrt{n}/5$ for each $i \in \{1, 2, \dots, n\}$. Since \mathcal{C}_k fluctuates by at most $a\sqrt{n}/5$ this also implies that the total number of burgers in $X(1, i)$ is at most $3a\sqrt{n}/5$ for each $i \in \{1, 2, \dots, n\}$, and hence $|X(1, i)| < a\sqrt{n}$ for each $i \in \{1, 2, \dots, n\}$. The result thus follows from Lemma 3.12 and the analogous bound that applies when \mathcal{D} is replaced by \mathcal{C} . \square

3.5. *Proof of Lemma 3.2.* Lemma 3.2 should not seem very surprising in light of the results we have established so far. Lemma 3.5 and Lemma 3.7 show that when the Lemma 3.2 hypothesis holds (i.e., $\chi \neq 2$) the stack $X(-\infty, 0)$, embedded in the manner of Figure 2, scales to a vertical line a.s. as one zooms out. It is thus natural to expect that the left-right fluctuation of the time-evolution of the stack is small compared to the up-down fluctuation. We divide Lemma 3.2 into Lemmas 3.14 and 3.16, which we state and prove below.

LEMMA 3.14. *If $\chi \neq 2$, then $\text{Var}[\mathcal{D}(n)] = o(n)$.*

PROOF. First, we claim that if $\chi \neq 2$ then $n^{-1/2}\mathcal{D}_n$ tends to zero in probability. To see this, recall that the stacks $X(-\infty, 0)$ and $X(-\infty, n)$ agree in law, and that the collection of the top $a\sqrt{n}$ burgers (for any fixed a) is likely to contain a roughly even distribution of hamburgers and cheeseburgers, in the sense of Lemma 3.7. By Lemma 3.13, the probability that the two stacks agree except for the top $a\sqrt{n}$ burgers [i.e., that one can find $j, k \leq a\sqrt{n}$ such that $X(-\infty, 0)$ with the top j burgers removed is the same as $X(-\infty, n)$ with the top k burgers removed] is a quantity that remains bounded below as n tends to infinity, and the bound can be made arbitrarily close to 1 by making a large enough. It follows from this that the random variables $n^{-1/2}\mathcal{D}_n$ tend to zero in probability. We still need to show that the variances $\mathbb{E}[n^{-1}\mathcal{D}_n^2]$ tend to zero. This follows from the fact that the random variables $n^{-1}\mathcal{D}_n^2$ tend to zero in probability, together with the uniform bounds on their tails given by Lemma 3.12. \square

LEMMA 3.15. *If $\text{Var}[\mathcal{D}(n)] = o(n)$ then $n^{-1/2} \max\{|\mathcal{D}_j| : 1 \leq j < nt\}$ converges to zero in probability for each fixed $t > 0$.*

PROOF. The variance assumption immediately implies that $n^{-1/2}\mathcal{D}_{\lfloor nt \rfloor}$ converges to zero in probability for each fixed t . It also implies that the joint law of $n^{-1/2}\mathcal{D}_{\lfloor nt \rfloor}$ at any fixed collection of t values tends to zero in probability.

Now, suppose that we divide the first t units of time into equal increments of length δ (for some small δ). Then on each such increment, we would guess that with high probability the fluctuation of $n^{-1/2}\mathcal{D}_{\lfloor tn \rfloor}$, as t ranges through an interval of length δ , should be very small when n is sufficiently large and δ is small and fixed. Indeed, the bound in Lemma 3.12 (applied with $a = \delta^{1/6}$) implies that the probability that there is even a single interval on which the fluctuation magnitude is greater than $a\delta^{1/2} = \delta^{2/3}$ will remain bounded above by some constant as $n \rightarrow \infty$, and this constant can be made arbitrarily small by taking δ sufficiently small. Since we can take δ as close to zero as we like, the lemma follows. \square

LEMMA 3.16. *If $\chi \neq 2$ and $\text{Var}[\mathcal{D}(n)] = o(n)$, then*

$$\lim_{n \rightarrow \infty} \mathbb{E}[\mathcal{D}(-n, -1) | 1_{J > n}] = 0.$$

PROOF. Let us assume $\text{Var}[\mathcal{D}(n)] = o(n)$ and proceed to derive the conclusion of the lemma. Let I be the smallest value of $j \geq 0$ for which $\mathcal{C}(-j, -1) = 1$. For each n , let μ_n be the measure whose Radon–Nikodym derivative with respect to μ is given by the (a.s. nonnegative) quantity $(1 - \mathcal{C}(-n, -1))1_{I \geq n}$. Since $(1 - \mathcal{C}(-n, -1))$ is a martingale (which reaches zero at time I) the optional stopping theorem implies that the expectation of $(1 - \mathcal{C}(-n, -1))1_{I \geq n}$ is 1, so that μ_n is in fact a probability measure. Informally, μ_n is the measure obtained from μ by conditioning on $a_j := (1 - \mathcal{C}(-j, -1))$ being positive until time n and multiplying the probability of each $X(-n) \cdots X(-1)$ sequence by a quantity proportional to a_n .

Note that if we just consider the law μ conditioned on having a_j eventually hit some $C > n$ before hitting zero, then the law of $\{a_1, a_2, \dots, a_n\}$ under this conditional law agrees with μ_n . In both measures, if we are given the values of a_i up to some $j \in \{0, 1, \dots, n - 1\}$, then the conditional probabilities that a_j goes up to $a_j + 1$ and down to $a_j - 1$ are proportional, respectively, to $a_j + 1$ and $a_j - 1$. Thus, there exists a measure μ_∞ whose law restricted to $X(j) : j \leq n$ agrees with μ_n , for each n . One can sample from μ_∞ by first sampling the a_n for all n (an ordinary simple random walk for negative n , a walk “conditioned to stay positive for all time” for positive n), which determines $\mathcal{C}(-n, -1)$, and then conditioned on that, choosing the burger and order types for each step independently from the usual conditional laws. It is well known that the μ_∞ law of the process $n^{-1/2}a_{\lfloor tn \rfloor}$ (where n is fixed, t is a parameter) converges to that of a three-dimensional Bessel process as n tends to infinity (which can be understood as “Brownian motion conditioned to stay positive for all time”). In particular, this implies that for every $\delta > 0$ there exists a $b = b(\delta) > 0$ such that

$$(23) \quad \mu_n \{ b^{-1} < n^{-1/2} |\mathcal{C}(-n, -1)| < b \} \geq 1 - \delta$$

for all sufficiently large n .

Using the measure μ_∞ allows us to convert an expectation into a probability: by definition, since $J \leq I$,

$$(24) \quad \mathbb{E}_\mu[|1 - \mathcal{C}(-n, -1)|1_{J>n}] = \mu_n\{J > n\} = \mu_\infty\{J > n\}.$$

We remark that by (2), (3), (4) and our assumption that $\chi \neq 2$, together with the fact that the quantity $\mathcal{C}(-n, -1)1_{J>n}$ that appears in (4) is nonpositive, the values $\mu_\infty\{J > n\}$ converge to a positive constant as $n \rightarrow \infty$. In other words, there is a positive μ_∞ probability that $J = \infty$, even though J is μ a.s. finite. An analog to (24) is the following:

$$(25) \quad \mathbb{E}_\mu[|\mathcal{D}(-n, -1)|1_{J>n}] = \mathbb{E}_{\mu_\infty}\left[\frac{|\mathcal{D}(-n, -1)|}{1 - \mathcal{C}(-n, -1)}1_{J>n}\right].$$

The quantity on the right is an expectation of a value that is bounded above. Hence, to show that it tends to zero, it suffices to show that the quantity $\frac{|\mathcal{D}(-n, -1)|}{1 - \mathcal{C}(-n, -1)}$ tends to zero in probability under the μ_∞ measure.

For this, in light of (23), it suffices to show that for every $b > 0$ and $b' > 0$, we have

$$(26) \quad \lim_{n \rightarrow \infty} \mu_\infty\{b^{-1} < n^{-1/2}|\mathcal{C}(-n, -1)| < b \text{ and } n^{-1/2}|\mathcal{D}(-n, -1)| > b'\} = 0.$$

We now fix such a b and b' and proceed to prove (26). First, let us make a remark about the strategy. We want to understand the measure μ_∞ (where we condition a_j to be positive) but most of the results in this paper apply to μ (which does not have such conditioning). The rough idea that helps us make the connection is that if we produce the simple walk a_j using μ and then *recenter* at a place where a_j is locally minimal, then the recentered process looks (locally) similar to a sample from μ_∞ .

Precisely, we may sample from μ and let \tilde{j} be the (smallest) value $j \in \{0, 1, \dots, 2n\}$ where a_j obtains its minimum. With probability at least 0.25, we have $\tilde{j} \in \{0, 1, \dots, n/2\}$ for even n . Conditioned on j and $a_{\tilde{j}}$, the law of a_k for $k \geq j$ is just that of a simple random walk conditioned not to go below a_j until time $2n$. We claim that the conditional μ law of $a_{j-\tilde{j}} - a_{\tilde{j}}$ for $j \in \{1, 2, \dots, n\}$ is similar to the μ_∞ law of a_j for $j \in \{1, 2, \dots, n\}$ in the sense that the Radon–Nikodym derivative between the two remains bounded between positive constants as n tends to infinity.

Indeed, we can describe this Radon–Nikodym derivative explicitly. We have two probability measures on sequences of the form r_1, r_2, \dots, r_n . The Radon–Nikodym derivative is a function of r_n alone; recall that we restricting attention to paths for which $r_n < 2bn^{1/2}$. In the μ_∞ case, the probability of each path is proportional to r_n , which is in turn proportional to the probability that a random walk started at r_n reaches height $2b\sqrt{n}$ before reaching zero. In the case of the conditional μ law, the probability of a path is proportional to the probability that walk started at

r_n fails to hit zero until at least the end what was the original length $2n$ interval (i.e., at least $n/2$ more steps). It is not hard to see that given that a particular one of these two events occurs (reaching height $2b\sqrt{n}$ before zero or staying positive for $n/2$ more steps) the conditional probability that the other occurs is at least some constant (which does not tend to zero as $n \rightarrow \infty$), and that this implies the claim.

In light of the above claim, (26) follows from Lemma 3.15. If the μ_∞ law of $n^{-1/2}\mathcal{D}(-n, -1)$ [conditioned on $b^{-1} < n^{-1/2}|\mathcal{C}(-n, -1)| < b$] failed to converge to zero, then the μ law of the maximum of $(2n)^{-1/2}|\mathcal{D}(-j, -1)|$ obtained on $(0, 2n)$ would have to also fail to converge to zero uniformly, contradicting Lemma 3.15. \square

3.6. *Proof of Theorem 2.5.* We now complete the proof of Theorem 2.5. The case $p \geq 1/2$ follows from Lemma 3.1 and Lemma 3.15, so we may assume $p < 1/2$. For each fixed value of t , the variance limits described in Lemma 3.1 guarantee that the variance of $n^{-1/2}\mathcal{D}_{\lfloor tn \rfloor}$ converges to αt as $n \rightarrow \infty$, where $\alpha = \max\{1 - 2p, 0\}$ as in the statement of Theorem 2.5. This implies that, at least subsequentially, the random variables $n^{-1/2}A_{tn}$ converge in law to a limit as $n \rightarrow \infty$ for each fixed t . The same is true of the joint law of $n^{-1/2}A_{\lfloor tn \rfloor}$, where t ranges over a finite set of values t_1, t_2, \dots, t_k . Our first step toward proving Theorem 2.5 will be to show that for any such t_1, t_2, \dots, t_k this joint law converges in law to the law of the corresponding Brownian motion restricted to these values.

To begin, we claim that if $p < 1/2$ then $\mathbb{E}[|E|] = \infty$. We argue the contrapositive: that if $\mathbb{E}[|E|] < \infty$ then $p \geq 1/2$. Lemma 3.7 states that if $\mathbb{E}[|E|] < \infty$ then the fraction of $\boxed{\text{H}}$ symbols among the top n elements in $X(-\infty, 0)$ a.s. tends to 0.5 as $n \rightarrow \infty$. One may deduce from this [and the fact that, by stationarity, $X(-\infty, 0)$ agrees in law with $X(-\infty, \lfloor tn \rfloor)$] that $n^{-1/2}\mathcal{D}_{\lfloor tn \rfloor}$ must converge in probability to zero in this case; together with the bounds in Lemma 3.13 this implies that $\text{Var}[\mathcal{D}_n] = o(n)$ in this case, so that (by Lemma 3.1) we must have $p \geq 1/2$.

Now, by Lemma 3.7, we must have (since we are assuming $p < 1/2$) that as $n \rightarrow \infty$ the fraction of $\boxed{\text{F}}$ symbols among the leftmost n elements of $X(1, \infty)$ tends a.s. to zero. Thus, the number of $\boxed{\text{F}}$ symbols in $X(1, \lfloor tn \rfloor)$ is $o(n^{1/2})$ with probability tending to 1 as $n \rightarrow \infty$. Thus if $t_3 = t_1 + t_2$ then $A_{\lfloor t_3 n \rfloor}$ is equal to $A_{\lfloor t_1 n \rfloor}$ plus something with the law of $A_{\lfloor t_2 n \rfloor}$ plus an error which is $o(n^{1/2})$ [which arises when we determine the values of Y_j for the X_j in the sequence that are equal to $\boxed{\text{F}}$, and also take into account the $O(1)$ rounding error from the $\lfloor \cdot \rfloor$ with high probability. Similarly, $A_{\lfloor tn \rfloor}$ is equal in law to the sum of k independent copies of $A_{\lfloor (t/k)n \rfloor}$ plus a term that is $o(n^{1/2})$ with high probability. From this it follows that any subsequential weak limit of the random variable $n^{-1/2}A_{\lfloor tn \rfloor}$ is infinitely divisible with the appropriate exponent and mean zero, hence a centered Gaussian with some covariance matrix; moreover, if we consider a finite collection of t values, the corresponding limiting joint law has independent Gaussian increments. Next, we claim that $\mathcal{C}_{\lfloor (t/k)n \rfloor}$ and $\mathcal{D}_{\lfloor (t/k)n \rfloor}$ converge separately to Brownian motions each with the correct variance. Indeed, recalling the variance limit described

in Lemma 3.1, this follows from the tightness of the random variable sequences $n^{-1}\mathcal{D}_n^2$ and $n^{-1}\mathcal{C}_n^2$, which is implied by the uniform decay bound of Lemma 3.13. This determines the diagonal elements of the diffusion covariance matrix for the limit of $n^{-1/2}A_{\lfloor tn \rfloor}$, but we still need to rule out the possibility of off-diagonal elements. However, it suffices to observe that the law must be symmetric under the operation that replaces all $\{\textcircled{\text{H}}, \boxed{\text{H}}\}$ symbols with the corresponding $\{\textcircled{\text{C}}, \boxed{\text{C}}\}$ symbols, so that by symmetry the limiting covariance between the \mathcal{C} and \mathcal{D} components must be zero. The extension from the finite-dimensional convergence result above to the stronger form of convergence claimed in the theorem statement follows from exactly same argument used to establish the $p \geq 1/2$ case in the proof of Lemma 3.15 (noting that Lemma 3.12 can be used to bound the oscillation of \mathcal{D} on small intervals).

REMARK 3.17. When $p \leq 1/2$, the conclusion of Theorem 2.5 holds for $t \in [0, \infty)$ even if we condition on an arbitrary time-zero burger stack S_0 in place of $X(-\infty, 0)$, because the number of $\boxed{\text{F}}$ symbols in $X(1, n)$ is $o(n^{1/2})$ with high probability. When $p > 1/2$, the argument combined with the monotonicity results shows that the convergence still holds if we condition on an initial stack that is *well balanced* in the sense that the fraction of hamburgers among its top k elements tends to $1/2$ as k tends to infinity.

4. Random planar maps.

4.1. *Bijection.* We begin by recalling a few classical constructions. A *planar map* M is a planar graph together with a topological embedding into the sphere (which we represent as the compactified complex plane $\mathbb{C} \cup \{\infty\}$). Self loops and edges with multiplicity are allowed. Some of the first enumeration formulas for planar maps were given by Mullin and by Tutte in the 1960s [47, 57] and since then a sizable literature on enumerations of planar maps of various types (and various bijections with labeled trees, walks, pairs of trees, etc.) has emerged. We describe only the very simplest formulations here. Let $V = V(M)$ be the set of vertices of M and $F = F(M)$ the set of faces. Let $Q = Q(M)$ be the map whose vertex set is $V \cup F$, and whose edge set is such that each $f \in F$ is connected to all the vertices along its boundary (see Figure 6). In other words, Q is obtained from M by adding a vertex to the center of each face and then joining each such vertex to all of the vertices (counted with multiplicity) that one encounters while tracing the boundary of that face. In particular Q is bipartite, with the two partite classes indexed by V and F , and all of the faces of Q are quadrilaterals (with one quadrilateral for each edge of M). Let M' denote the dual map of M (the map whose vertices correspond to the faces of M —an edge joins two vertices in M' if an edge borders the corresponding faces in M). See Figure 6.

Now, suppose that M is endowed with a distinguished oriented edge e . This determines an oriented edge e_0 of Q that has the same initial endpoint as e and is

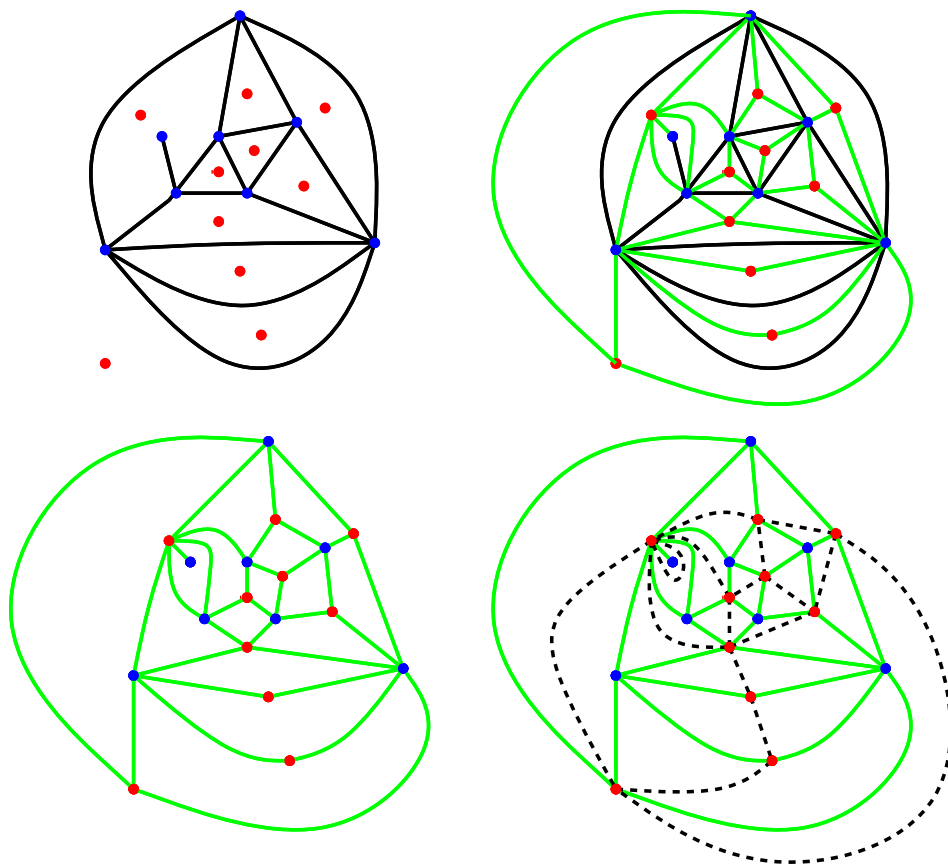


FIG. 6. Upper left: a planar map M with vertices in blue and “dual vertices” (one for each face) shown in red. Upper right: the quadrangulation $Q = Q(M)$ formed by adding a green edge joining each red vertex to each of the boundary vertices of the corresponding face. Lower left: quadrangulation Q shown without M . Lower right: the dual map M' corresponding to the same quadrangulation, obtained by replacing the blue-to-blue edge in each quadrilateral with the opposite (red-to-red) diagonal.

the next edge clockwise (among all edges of M and Q that start at that endpoint) from e_0 . We refer to the endpoint of e_0 in V as the *root* and the endpoint in F as the *dual root*. Now, suppose that T is a subset of the edges of M corresponding to a spanning tree. Then the set T' of dual edges to the edges in the complement of T is necessarily a spanning tree of M' , see Figure 7. The union of T , T' and Q forms a triangulation, with each triangle containing two edges from Q and one from either T or T' . Let $e_0, e_1, e_2, \dots, e_{2n} = e_0$ be the sequence of edges of this triangulation hit by the path shown in Figure 7, which starts at a midpoint of e_0 and then crosses each edge of Q exactly once (with an element of V on the left and an element of F on the right) before returning to the initial edge. This path goes




FIG. 7. *Left: in each quadrilateral we either draw an edge (connecting blue to blue) or the corresponding dual edge (connecting red to red). In this example, the edges drawn form a spanning tree of the original (blue-vertex) graph, and hence the dual edges drawn form a spanning tree of the dual (red-vertex) graph. Right: designate a “root” (large blue dot) and an adjacent “dual root” (large red dot). The red path starts at the midpoint of the green edge between the root and the dual root and crosses each of the green edges once, keeping the blue endpoint to the left and red endpoint to the right, until it returns to the starting position.*

through each triangle without ever crossing an edge of T or T' (and in a sense it describes the boundary between T and T').

For each e_i , let $d(e_i) = (d_1, d_2)$, where d_1 is the number of edges in the tree T in between the V endpoint of e_i and the root, and e_2 is the number of edges in the tree T' in between the F endpoint of e_i and the dual root. Then the sequence $d(e_0), d(e_1), \dots, d(e_{2n}) = d(e_0)$ is a simple walk on the lattice \mathbb{Z}_+^2 of nonnegative integer pairs. We can associate a corresponding word in the alphabet Θ by writing $\textcircled{\text{H}}$ or $\textcircled{\text{C}}$ each time the first or second (respectively) coordinate of $d(e_i)$ goes up and an $\boxed{\text{H}}$ or $\boxed{\text{C}}$ each time the first or second (respectively) coordinate of $d(e_i)$ goes down. The following is the word corresponding to Figure 7:

$\textcircled{\text{C}}\textcircled{\text{C}}\textcircled{\text{C}}\textcircled{\text{H}}\textcircled{\text{C}}\textcircled{\text{H}}\textcircled{\text{H}}\boxed{\text{C}}\boxed{\text{H}}\boxed{\text{C}}\boxed{\text{C}}\boxed{\text{C}}\boxed{\text{C}}\textcircled{\text{C}}\textcircled{\text{C}}\textcircled{\text{H}}\textcircled{\text{H}}\boxed{\text{C}}\boxed{\text{H}}\boxed{\text{C}}\boxed{\text{C}}\boxed{\text{C}}\textcircled{\text{H}}\textcircled{\text{C}}\textcircled{\text{H}}\textcircled{\text{H}}$
 $\boxed{\text{C}}\boxed{\text{C}}\boxed{\text{H}}\textcircled{\text{H}}\textcircled{\text{H}}\textcircled{\text{H}}\boxed{\text{C}}.$

In fact, this construction describes a bijection (essentially due to Mullin [47] but more explicitly explained by Bernardi [6]) between the set of pairs (M, T) —where M is rooted planar map with n edges and T a distinguished spanning tree—and the set of walks of length $2n$ in \mathbb{Z}_+^2 that start and end at the origin. This set of walks is in turn equivalent to the set of length $2n$ words W in the symbols $\{\textcircled{\text{C}}, \textcircled{\text{H}}, \boxed{\text{C}}, \boxed{\text{H}}\}$ for which $\overline{W} = \emptyset$.

We can say more about this bijection. Every quadrilateral of Q corresponds to a burger. The quadrilateral is divided by an edge in $T \cup T'$ into two triangles. The

first triangle the path goes through corresponds to the step at which that burger was added to the stack, while the second corresponds to the step at which the same burger was ordered. Quadrilaterals of Q that are divided by T edges correspond to hamburgers while elements divided by T' edges correspond to cheeseburgers. Another equivalent point of view is that every vertex of T (besides the root) corresponds to a hamburger and every vertex of T' (besides the dual root) corresponds to a cheeseburger (see Figure 8). Let U_k be the union of the first k triangles traversed. Then a burger is added the first time the corresponding vertex is part of U_k , and ordered the first time the corresponding vertex lies in the interior of U_k . The outer boundary vertices of U_k (excluding the root and the dual root) represent the

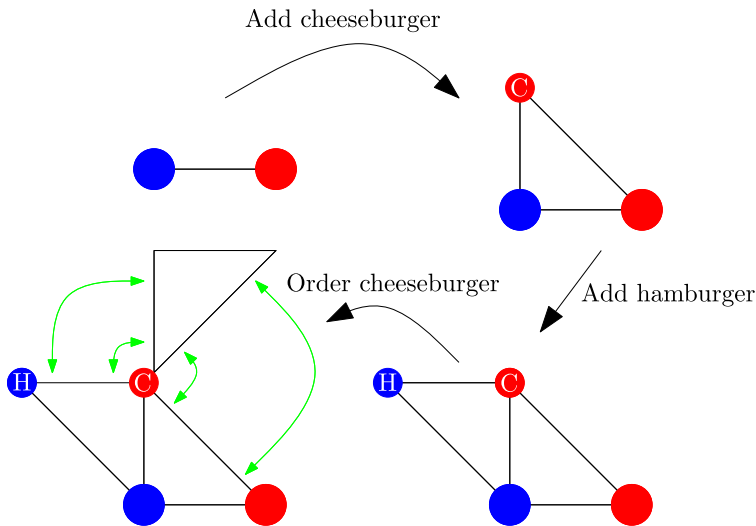


FIG. 8. Start with a single edge (upper left). We will recursively construct a triangulation with two types of vertices by “exploring” triangles one at a time (cf. Angel’s peeling process [1]). At each step the outer boundary of the explored region has two edges connecting two vertices of different types: one “special edge” (to which a new triangle will be glued) and one “root edge” (the initial edge in the construction; the special edge and root edge coincide at the beginning). When a new triangle is added, there are four possibilities for its “type,” corresponding to the four symbols. Adding a triangle corresponding to one of the two “burger” symbols amounts to drawing a correspondingly colored vertex outside the surface and making a triangle by connecting it to the endpoints of the special edge (i.e., the clockwise-most boundary hamburger and the counterclockwise-most boundary cheeseburger). Adding a triangle corresponding to one of the “order” symbols amounts to making a new triangle by connecting one endpoint of the special edge to another vertex adjacent to the opposite endpoint of the special edge. At any given time in this process, the small dots on the outer boundary of the explored region correspond to burgers that have not yet been consumed. To put this somewhat fancifully, an illiterate restaurant owner armed only with triangular pieces of paper and cheeseburger and hamburger stickers could use this scheme to keep track of the day’s events, gradually constructing a surface as the day goes on.




FIG. 9. *Left: The subset of edges does not form a spanning tree, so we obtain multiple red loops (again, each green line is crossed once). Right: a canonical way to replace some edges with dual edges, and vice versa, to obtain a spanning tree.*

burgers on the stack at time k . (The fact that T and T' are trees ensures that U_k will be simply connected for all k .)

Now what happens if we remove the requirement that T be a tree? Let T be any subset of the edges of M and let T' be the set of dual edges of the edges in the complement of T . As shown in the left side of Figure 9, we still have a collection of loops, each of which passes through some subset of the triangles (crossing only edges of Q , no edges of T or T') such that every edge in Q is crossed exactly once.

Observe carefully the loops on the left-hand side of Figure 9. We will now describe a way (corresponding to the right-hand side of Figure 9) to modify the pair T, T' (replacing some edges with dual edges, and vice versa) in a canonical way to obtain a tree/dual-tree pair \tilde{T}, \tilde{T}' .

Let L_0, L_1, \dots be the loops (as shown in Figure 9), where L_0 is the special loop that passes through the edge e_0 . Each loop can also be identified with the set of triangles it passes through and viewed as a subset of the set \mathcal{T} of triangles in $T \cup T' \cup Q$. Note that \mathcal{T} can itself be viewed as a graph (two triangles adjacent if they share an edge). Let C_1, \dots, C_k be the components of $\mathcal{T} \setminus L_0$. (If we remove from the surface all of the triangles that L_0 passes through, then each C_j is a connected subset of the triangles that remain.) Clearly, L_0 passes through at least one triangle on the outer boundary of each C_j (i.e., at least one triangle that shares an edge with a triangle in C_j). Let A_j describe the *last* triangle sharing an edge with the boundary of C_j that is traversed by L_0 .

This edge is either an edge of T or an edge of T' . In either case, we will now *replace it* with the opposite diagonal of the same quadrilateral (i.e., we replace a dual edge with an edge or an edge with a dual edge). The effect of all of these

replacements is to join one loop in each of the C_j to the primary loop L_0 . The total number of loops has decreased by k . We then repeat this process until at the end there is only a single loop dividing a tree from a dual tree, which we denote by \tilde{T} and \tilde{T}' . See Figure 9. In that figure we colored all of the edges in $\tilde{T} \setminus T$ and $\tilde{T}' \setminus T'$ a different color—we refer to these as “fictional edges” since they do not belong to the original pair T, T' .

To describe this (M, T) by a word in Θ , we first construct the word \tilde{W} in $\{\textcircled{C}, \textcircled{H}, \boxed{C}, \boxed{H}\}$ corresponding to \tilde{T} and \tilde{T}' via the method described above. We then make an observation about \tilde{W} . We already know from above that whenever the second triangle of a quadrilateral is traversed by the path corresponding to \tilde{T} and \tilde{T}' , it corresponds to an order \boxed{C} or \boxed{H} . We claim that whenever the second triangle of a “fictional” quadrilateral is traversed, it corresponds to the consumption of a burger on the top of the stack. To see this, let C_1, \dots, C_k be as above and consider the boundary edges of one of the C_j . Since these boundary edges cannot include any edges of Q , they consist only of edges in T or T' . Since the boundary of C_j is connected, it contains either only edges of T (and vertices of V) or only edges of T' (and vertices of F). Now, there is one edge or dual edge on the boundary of C_j that is altered—replaced by a fictional edge or dual edge—to create two new triangles. The reader may observe that in between the time that the final path goes through the first of these triangles and the time it goes through the second triangle, it does not traverse the first half of any other quadrilateral without traversing also the second half. In other words, no burger added after the first triangle is traversed remains on the stack when the second triangle is traversed.

Now to define the word W that corresponds to (M, T) , we start with \tilde{W} and then put an \boxed{F} in place of each \boxed{C} or \boxed{H} that corresponds to a triangle with a fictional edge or dual edge. For example, the word corresponding to Figure 9 is

$$\textcircled{C}\textcircled{H}\boxed{C}\textcircled{H}\textcircled{C}\textcircled{C}\textcircled{C}\textcircled{H}\textcircled{C}\textcircled{C}\textcircled{C}\textcircled{H}\boxed{C}\textcircled{H}\textcircled{H}\boxed{C}\textcircled{H}\textcircled{C}\textcircled{C}\textcircled{H}\boxed{C}\textcircled{C}\textcircled{C}$$

$$\boxed{H}\boxed{H}\boxed{F}\boxed{H}\boxed{F}\textcircled{H}\boxed{F}\boxed{F}\boxed{F}.$$

As noted above, \tilde{W} corresponds to the same burger production/consumption sequence as W . Given a sequence of this type, it is straightforward to reconstruct the corresponding planar map. Namely, we first replace each \boxed{F} with the corresponding \boxed{C} or \boxed{H} , then construct the corresponding map with \tilde{T} and \tilde{T}' , and then reverse the edges and dual edges corresponding to \boxed{F} symbols to create T and T' . Essentially the same argument as above shows that we recover the same \tilde{T} and \tilde{T}' when we follow the algorithm to create them from T and T' . We have now given a bijection between length $2n$ words W in Θ with $\overline{W} = 0$ and pairs (M, T) , where M is a rooted planar map with n edges and T is any distinguished subset of the edges of M . In the case that M is a fixed planar map, this is equivalent to a construction given by Bernardi in Section 4 of [8], as alluded to in the Introduction.

TABLE 2
Correspondences between decorated map and word

Rooted map/edge set: (M, T)	Corresponding word: W
Number of edges in \tilde{T}	Number of \textcircled{H} 's in W
Number of edges in \tilde{T}'	Number of \textcircled{C} 's in W
Number of loops	1 plus number of \boxed{F} 's in W
Number of components of T	1 plus number of \boxed{F} 's matched to \textcircled{C} 's
Number of components of T'	1 plus number of \boxed{F} 's matched to \textcircled{H} 's
Number of edges in T	Number of \textcircled{H} to \boxed{H} or \textcircled{C} to \boxed{F} matches
Number of edges in T'	Number of \textcircled{C} to \boxed{C} or \textcircled{H} to \boxed{F} matches

Specifically, Section 4 of [8] describes a mechanism for constructing \tilde{T} from T that is equivalent to the one described above.

Each loop $L \neq L_0$ has an inside (the component of the sphere minus L that contains L_0) and an outside (the component containing L_0). There are two types of loops L : those that pass only through triangles with edges on the inside (dual edges on the outside) and those that pass only through triangles containing edges on the outside (dual edges on the inside). In other words, each loop can be viewed as a boundary interface between a cluster and a dual cluster, and the cluster can be either inside or outside of the loop. Since every cluster (dual cluster) has a unique loop tracing its outer boundary, we can count the clusters (dual clusters) by counting the number of loops of the two types, which in turn corresponds to counting the number of \boxed{F} 's matched to \textcircled{C} 's or to \textcircled{H} 's, as explained in Table 2.

Part of our motivation for introducing the bijection of this section is that a complicated “nonlocal” quantity on the planar map side (number of loops) becomes a more straightforward “local” quantity on the Θ -word side (number of \boxed{F} symbols).

4.2. *Critical Fortuin–Kastelyn model and infinite volume limits.* Now, suppose we fix a p , choose an $X(n)$ sequence from μ , and then condition on $\overline{W} = 0$ where $W = X(1) \cdots X(2n)$. Using the bijection above, this gives us a random rooted planar map M decorated by a subgraph T . The probability of seeing a particular (M, T) with ℓ loops (which corresponds to ℓ of the \boxed{F} characters) is then proportional to

$$(27) \quad \left(\frac{p/2}{(1-p)/4} \right)^\ell = (2p/(1-p))^\ell.$$

If for some $q > 0$, we write $p = \frac{\sqrt{q}}{2+\sqrt{q}}$, so that p solves the equation $2p/(1-p) = \sqrt{q}$, then (27) becomes equivalent to $\sqrt{q}^\ell = q^{\ell/2}$.

It is natural to choose a random (M, T) pair from the uniform probability measure on such pairs (with n total edges in M) weighted by a quantity proportional to $q^{\ell/2}$. One reason this is natural is that (once we condition on M) the law of T is that of the self-dual Fortuin–Kastelyn model, which is in turn related to the critical q -state Potts model, see [4]. We will not say more about this here, but the survey [32] contains one clear account of this connection [as well as conjectures, due to Duplantier, Nienhuis, and others, relating these models to SLE and conformal loop ensembles with $q = 2 + 2 \cos(8\pi/\kappa)$ when $q \in (0, 4)$ and $\kappa \in (4, 8)$].

Next, we claim that

$$(28) \quad \lim_{n \rightarrow \infty} \frac{1}{n} \log \mu\{X(1, n) = \emptyset\} = 0.$$

In other words, the probability that $X(1, n) = \emptyset$ tends to zero more slowly than any exponential function of n . (In fact, Gwynne and Sun [29] have recently proved that it decays like a negative power of n , and have derived the exact exponent as a function of p , but we will not need this here.)

Let $m = \lfloor \sqrt{n} \rfloor$ and note that it follows from Theorem 2.5 (using the type of argument in Lemma 3.13) that $X(m, n - m)$ has a probability bounded below (independently of n) of containing at most m elements. Conditioned on this, there is some positive probability that the values of $X(1), \dots, X(m)$ and $X(n - m + 1) \cdots X(n)$ are such that $X(1, n) = \emptyset$. This conditional probability cannot be smaller than exponentially small in m , and since $m = o(n)$, this implies (28).

Now we recall Cramér’s theorem (see, e.g., the reference text [16]). Suppose that Z_1, Z_2, \dots are i.i.d. random variables each of which takes values in $\{1, 2, \dots, k\}$ with resp. probabilities a_1, a_2, \dots, a_k . Write $a = (a_1, a_2, \dots, a_k)$ for the probability vector. Let u^n be a vector encoding the empirical distribution of Z_1, Z_2, \dots, Z_n . That is, the j th component of the vector u^n , which we may denote by u^n_j , is the fraction of the n elements Z_1, Z_2, \dots, Z_n that are equal to j . Then Cramér’s theorem states that if A is any neighborhood in \mathbb{R}^k about the point a , then $\mathbb{P}\{u^n \notin A\}$ decays exponentially as $n \rightarrow \infty$.

Now, fix some large value for m and let $Z_j = (X_{jm}, X_{jm+1}, X_{jm+2}, \dots, X_{j(m+m-1)})$. That is, the Z_j correspond to length- m blocks of elements from the sequence X . Now, if the X_1, X_2, \dots are i.i.d. then the corresponding Z_1, Z_2, \dots are i.i.d. and Cramér’s theorem implies that $\mathbb{P}\{u^n \notin A\}$ decays exponentially as $n \rightarrow \infty$. But (28) implies that the $\mathbb{P}\{X(1, mn + r) = \emptyset\}$ decays *subexponentially* in n (uniformly over values of $r \in \{0, 1, 2, \dots, m - 1\}$). This implies that if we condition on $X(1, mn + r) = \emptyset$ then (for any open set $A \subset \mathbb{R}^k$, where k is the number of possible values Z_1 as defined above can assume, and A contains the probability vector whose ℓ th entry is the probability that Z_1 assumes the ℓ th of its possible values) the conditional probability that $u^n \notin A$ tends to zero as $n \rightarrow \infty$ (uniformly over choices for r).

From this, we may conclude that if one chooses $X(1) \cdots X(N)$ conditioned on $X(1, N) = \emptyset$ and then picks a random k and recenters the sequence at $X(k)$, then one obtains a sequence of random recentered processes whose infinite volume limit is a random process with law μ . In this sense, μ describes an infinite volume limit of the critical Fortuin–Kastelyn models. Given a sequence $X(\cdot)$ sampled from μ , the corresponding sequence $Y(\cdot)$ can be then used to construct an infinite random surface with an infinite spanning tree and spanning dual tree (in the manner described above for finite words). From this perspective, Theorem 2.5 can be understood a scaling limit result about this random pair of infinite trees (which in turn encode the structure of a random infinite planar map).

APPENDIX: BACKGROUND: CONTEXT AND MOTIVATION

We now say a bit more about the larger project that motivated this work, namely, the problem of relating discrete random surfaces to continuum objects like Liouville quantum gravity, conformal loop ensembles, and the Schramm–Loewner evolution.

A longstanding open physics-motivated problem is to show that as the number of edges in the random maps described here tends to infinity, the corresponding loop-decorated surfaces (appropriately rescaled) converge in law to a limiting loop-decorated random surface \mathcal{M} . The physics literature provides ample heuristic evidence, in various settings, for the existence of a phase transition analogous to the $p = 1/2$ transition we present: $p > 1/2$ corresponds to what is called the *branched polymer* (a.k.a. *continuum random tree*) phase, while $p < 1/2$ is the *Liouville quantum gravity* phase. The literature is too vast for us to properly survey here, but sample works in this direction include [3, 9–11, 13–15, 17, 18, 23, 24, 31, 34, 35, 56]. This paper presents a clear and rigorous illustration of the phase transition.

The $p > 1/2$ case is relatively simple. Applying Theorem 2.5 when $p > 1/2$ we see that in the scaling limit, the spanning tree and the dual spanning tree that we construct converge in law to a.s. identical continuum random trees. The collection of quadrilaterals corresponding to the tree and the collection corresponding to the dual tree can each be interpreted as approximations of the same continuum random tree: gluing the two together produces a surface that, as a metric space, should approximate the same tree. [Note that since $X(-\infty, n)$ corresponds to a branch of a tree and $X(n, \infty)$ a branch of the corresponding dual tree, Proposition 3.8 implies that when we glue these surfaces together a positive fraction of the points along any branch of the tree or dual tree correspond to length-one loops that are “bottlenecks” of the combined surface. We do not know whether this remains true when $p = 1/2$.]

The $p < 1/2$ case is more interesting. Based on various theorems and heuristics, we expect the limiting surface \mathcal{M} to be a *Liouville quantum gravity* (LQG) surface decorated by a *conformal loop ensemble* (CLE), with respective parameters γ and

κ depending on q . Precise formulations of this conjecture, along with definitions of LQG surfaces (whose laws depend on a parameter γ) and CLEs (whose laws depend on a parameter κ , and are based on the Schramm–Loewner evolution with parameter κ), can be found in [21, 22, 52, 53] (along with much more extensive lists of references; the CLEs relevant to this paper have a parameter $\kappa \in (4, 8]$, while the $\kappa \leq 4$ case is discussed in [54]). But in a sense this question has been around since the 1980s when Polyakov first introduced Liouville quantum gravity [48, 49] (see also [33]).

Polyakov writes in a recent memoir [50] that he first became convinced of the connection between the discrete models and Liouville quantum gravity after jointly deriving, with Knizhnik and Zamolodchikov, the *KPZ formula* for so-called LQG scaling dimensions and comparing these dimensions with known combinatorial results for the discrete models [33]. The KPZ formula was recently proved mathematically in [20, 21] and can indeed be interpreted as evidence for the convergence conjectures. But actually proving these conjectures, even in a physics sense, remains a challenge.

There are various ways to formulate a convergence statement, depending on what topology one uses when talking about convergence in law [52]. One approach is to consider a model without loops and to interpret the discrete surfaces as random metric spaces. If one uses the Gromov–Hausdorff topology on metric spaces, then recent independent works by Miermont and Le Gall show that these metric spaces converge in law to a limiting random metric space called the *Brownian map* [39, 44]. Prior to these recent works, it was known that these discrete random metric spaces converge in law to limiting random metric spaces along subsequences, and that any subsequential limit is a random metric space that is almost surely homeomorphic to the sphere; see [37, 38, 41, 43], as well as the recent survey [40]. The Brownian map is conjecturally equivalent, as a metric space, to an LQG surface with $\gamma = \sqrt{8/3}$ [52].⁴ So far, there has been little success in extending the metric space theory to more general FK-weighted random surfaces, which are believed to correspond to other values of γ .

Another approach to the convergence problem is to ignore metric space structure and focus instead on conformal structure. In this case, one may consider a conformal map from a random discretized surface to the sphere, and then study the induced area measure on the sphere and the images of the loops on the sphere. The conformal point of view plays a central role in the physics literature on Liouville quantum gravity and “random metrics”, which has been developed as part of a general subject called conformal field theory. (In this literature, the term “metric” refers most directly to a Riemannian metric tensor, not a two-point distance

⁴One barrier to establishing this equivalence is that it remains an open question even to show that LQG surfaces *have* a canonical metric space structure. There is an obvious way to define a two-point distance function on an LQG surface in terms of limits, as explained in a footnote in [52], but it has not been proved that these limits exist and are nontrivial.

function.) Various precise conjectures along these lines are presented in [21, 52], along with a more extensive bibliography.

This paper takes a third approach. We encode loop-decorated surfaces by walks on \mathbb{Z}^2 (“inventory trajectories”). Roughly speaking, the loops determine a canonical non-self-intersecting “exploration path” which goes through every face on the surface and more or less traces all the loops. This path divides the surface into a “tree” and a “dual tree”—and the exploration driving function encodes the structure of this pair of trees. If we consider a topology in which two loop-decorated surfaces are considered close when their exploration driving functions are close, then Theorem 2.5 can be interpreted as a scaling limit result for the random surfaces. It describes the limiting law of the pair of trees and the manner in which they are glued together. (Note: the symmetry between the tree and dual tree makes our approach very different from the version of the Schaeffer bijection [51] used to define the Brownian map [40]. The latter involves a breadth-first search tree with geodesic branches and a dual tree that looks completely different from the tree itself—in particular, its diameter scales like a different power of the number of edges.)

It turns out that if one replaces discrete loops with a continuum CLE, then the loops determine a continuum (space-filling) analog of the “exploration path” described above, which traces through all of the loops in a canonical way [45, 53]. We can interpret this continuum exploration process as tracing the boundary between a continuum (space-filling) tree and a continuum (space-filling) dual tree. If one draws these trees on a Liouville quantum gravity random surface, then the lengths of branches of the tree are well defined a.s. by the results of [52], so that one can construct a continuum analog of the exploration driving function described in this paper.

As mentioned in the [Introduction](#), since this paper was originally posted the arXiv there have been several additional papers that have extended the theory developed here. One recently completed work with Bertrand Duplantier and Jason Miller [19] demonstrates that LQG surfaces decorated by CLEs indeed have well-defined exploration driving functions themselves, and that their laws are those of precisely the same kinds of Brownian motions as long as the parameters p, q, κ, γ are matched up correctly.⁵ As in the discrete case, the continuum exploration driving function should encode the metric space structure of a certain tree (here a continuum random tree) and a dual-tree. The simplest case is when $\kappa = 8$, corresponding to the case $p = 0$. The results of this paper are trivial (following from the classical central limit theorem for simple random walks) in that case, but the Liouville quantum gravity construction is still interesting.

⁵Specifically, $p \in [0, 1/2)$, $q \in [0, 4)$, $\kappa \in (4, 8]$, $\gamma \in [\sqrt{2}, 2)$ and $q = 2 + 2 \cos \frac{8\pi}{\kappa}$, $\gamma = \sqrt{16/\kappa}$,
 $p = \frac{\sqrt{q}}{2 + \sqrt{q}}$.

As explained in [19], this program allows us to interpret the results presented here to mean that the discrete loop-decorated surfaces of this paper converge to CLE-decorated Liouville quantum gravity in a topology where two configurations are close if the corresponding driving functions are close. Extensions of this result to stronger topologies (which encode lengths and intersection locations of loop configurations) have been established in works by (various subsets of) Cheng Mao, Ewain Gwynne, Jason Miller, and Xin Sun [27–30]. See also the related results in [5, 12].

Moving from this kind of convergence to convergence in topologies that more directly encode the metric space and/or conformal structure of the random surfaces themselves appears difficult, but conceivably possible. One way to do this would be to show that if one couples a discrete loop-decorated surface with a CLE-decorated LQG surface in such a way that their driving functions are likely to be close, then it is likely that the two surfaces are also in some sense close as metric spaces, or as Riemannian manifolds with a conformal structure. We expect statements of this kind to be true, but proving them will require new ideas.

Since we would ultimately like to extend the results presented here to other topologies, we conclude with one relevant question. If two discrete loop-decorated surfaces \mathcal{M}_1 and \mathcal{M}_2 are conditioned to have exploration driving functions that are close in the L^∞ sense for all time, is it then the case that (with high probability) one can conformally map \mathcal{M}_1 to \mathcal{M}_2 in such a way that the image of the exploration path in \mathcal{M}_1 is close to the exploration path in \mathcal{M}_2 (at least on a fixed compact interval of time)? This question is already interesting in the case that $p = 0$. The results by Gill and Rohde in [26] about Brownian motion on these random surfaces, and the parabolicity of the infinite volume surfaces, constitute a promising step in this direction.

Acknowledgments. We thank Olivier Bernardi, Bertrand Duplantier, Ewain Gwynne, Nike Sun, Xin Sun, and Sam Watson for very helpful comments on drafts of this paper. We thank an anonymous referee for proposing many improvements that were incorporated into the text.

REFERENCES

- [1] ANGEL, O. (2003). Growth and percolation on the uniform infinite planar triangulation. *Geom. Funct. Anal.* **13** 935–974. [MR2024412](#)
- [2] AZUMA, K. (1967). Weighted sums of certain dependent random variables. *Tôhoku Math. J.* (2) **19** 357–367. [MR0221571](#)
- [3] BAILLIE, C. F. and JOHNSTON, D. A. (1992). Multiple Potts models coupled to two-dimensional quantum gravity. *Phys. Lett. B* **286** 44–52.
- [4] BAXTER, R. J., KELLAND, S. B. and WU, F. Y. (1976). Equivalence of the Potts model or Whitney polynomial with an ice-type model. *J. Phys. A: Math. Gen.* **9** 397.
- [5] BERESTYCKI, N., LASLIER, B. and RAY, G. (2015). Critical exponents on Fortuin–Kastelyn weighted planar maps. Preprint. Available at [arXiv:1502.00450](#).

- [6] BERNARDI, O. (2007). Bijective counting of tree-rooted maps and shuffles of parenthesis systems. *Electron. J. Combin.* **14** Research Paper 9, 36 pp. (electronic). [MR2285813](#)
- [7] BERNARDI, O. (2008). A characterization of the Tutte polynomial via combinatorial embeddings. *Ann. Comb.* **12** 139–153. [MR2428901](#)
- [8] BERNARDI, O. (2008). Tutte polynomial, subgraphs, orientations and sandpile model: New connections via embeddings. *Electron. J. Combin.* **15** Research Paper 109, 53. [MR2438581](#)
- [9] BERNARDI, O. and BOUSQUET-MÉLOU, M. (2011). Counting colored planar maps: Algebraicity results. *J. Combin. Theory Ser. B* **101** 315–377. [MR2802884](#)
- [10] BOROT, G. and EYNARD, B. (2011). Enumeration of maps with self-avoiding loops and the $O(n)$ model on random lattices of all topologies. *J. Stat. Mech. Theory Exp.* **1** P01010, 62. [MR2770626](#)
- [11] CATES, M. E. (1988). The Liouville field theory of random surfaces: When is the bosonic string a branched polymer? *EPL (Europhysics Letters)* **7** 719.
- [12] CHEN, L. (2015). Basic properties of the infinite critical-FK random map. Preprint. Available at [arXiv:1502.01013](#).
- [13] DAUL, J.-M. (1995). Q-States Potts model on a random planar lattice. Preprint. Available at [arXiv:hep-th/9502014](#).
- [14] DAVID, F. (1985). Planar diagrams, two-dimensional lattice gravity and surface models. *Nuclear Phys. B* **257** 45–58. [MR0788098](#)
- [15] DAVID, F. (1997). A scenario for the $c > 1$ barrier in non-critical bosonic strings. *Nuclear Phys. B* **487** 633–649. [MR1432821](#)
- [16] DEMBO, A. and ZEITOUNI, O. (2010). *Large Deviations Techniques and Applications. Stochastic Modelling and Applied Probability* **38**. Springer, Berlin. [MR2571413](#)
- [17] DUPLANTIER, B. and KOSTOV, I. (1988). Conformal spectra of polymers on a random surface. *Phys. Rev. Lett.* **61** 1433–1437. [MR0960093](#)
- [18] DUPLANTIER, B. and KOSTOV, I. K. (1990). Geometrical critical phenomena on a random surface of arbitrary genus. *Nuclear Phys. B* **340** 491–541. [MR1068092](#)
- [19] DUPLANTIER, B., MILLER, J. and SHEFFIELD, S. (2014). Liouville quantum gravity as a mating of trees. Preprint. Available at [arXiv:1409.7055](#).
- [20] DUPLANTIER, B. and SHEFFIELD, S. (2009). Duality and the Knizhnik–Polyakov–Zamolodchikov relation in Liouville quantum gravity. *Phys. Rev. Lett.* **102** 150603, 4. [MR2501276](#)
- [21] DUPLANTIER, B. and SHEFFIELD, S. (2011). Liouville quantum gravity and KPZ. *Invent. Math.* **185** 333–393. [MR2819163](#)
- [22] DUPLANTIER, B. and SHEFFIELD, S. (2011). Schramm–Loewner evolution and Liouville quantum gravity. *Phys. Rev. Lett.* **107** 131305.
- [23] EYNARD, B. (2011). Formal matrix integrals and combinatorics of maps. In *Random Matrices, Random Processes and Integrable Systems*. 415–442. Springer, New York. [MR2858440](#)
- [24] EYNARD, B. and BONNET, G. (1999). The Potts- q random matrix model: Loop equations, critical exponents, and rational case. *Phys. Lett. B* **463** 273–279. [MR1722272](#)
- [25] GEORGII, H.-O. (2011). *Gibbs Measures and Phase Transitions*, 2nd ed. *de Gruyter Studies in Mathematics* **9**. de Gruyter, Berlin. [MR2807681](#)
- [26] GILL, J. T. and ROHDE, S. (2013). On the Riemann surface type of random planar maps. *Rev. Mat. Iberoam.* **29** 1071–1090. [MR3090146](#)
- [27] GWYNNE, E., MAO, C. and SUN, X. (2015). Scaling limits for the critical Fortuin–Kasteleyn model on a random planar map I: Cone times. Preprint. Available at [arXiv:1502.00546](#).
- [28] GWYNNE, E. and MILLER, J. (2015). Convergence of the topology of critical Fortuin–Kasteleyn planar maps to that of CLE_κ on a Liouville quantum surface. In preparation.

- [29] GWYNNE, E. and SUN, X. (2015). Scaling limits for the critical Fortuin–Kastelyn model on a random planar map II: Local estimates and empty reduced word exponent. Preprint. Available at [arXiv:1505.03375](https://arxiv.org/abs/1505.03375).
- [30] GWYNNE, E. and SUN, X. (2015). Scaling limits for the critical Fortuin–Kastelyn model on a random planar map III: Finite volume case. In preparation.
- [31] HARRIS, G. (1994). Percolation on strings and the cover-up of the $c = 1$ disaster. *Nuclear Phys. B* **418** 278–304. [MR1281359](https://arxiv.org/abs/MR1281359)
- [32] KAGER, W. and NIENHUIS, B. (2004). A guide to stochastic Löwner evolution and its applications. *J. Stat. Phys.* **115** 1149–1229. [MR2065722](https://arxiv.org/abs/MR2065722)
- [33] KNIZHNIK, V. G., POLYAKOV, A. M. and ZAMOŁODCHIKOV, A. B. (1988). Fractal structure of 2D-quantum gravity. *Modern Phys. Lett. A* **3** 819–826. [MR0947880](https://arxiv.org/abs/MR0947880)
- [34] KOSTOV, I. K. (1989). $O(n)$ vector model on a planar random lattice: Spectrum of anomalous dimensions. *Modern Phys. Lett. A* **4** 217–226. [MR1016518](https://arxiv.org/abs/MR1016518)
- [35] KOSTOV, I. K. (1989). The ADE face models on a fluctuating planar lattice. *Nuclear Phys. B* **326** 583–612. [MR1028516](https://arxiv.org/abs/MR1028516)
- [36] LE GALL, J.-F. (1995). The Brownian snake and solutions of $\Delta u = u^2$ in a domain. *Probab. Theory Related Fields* **102** 393–432. [MR1339740](https://arxiv.org/abs/MR1339740)
- [37] LE GALL, J.-F. (2007). The topological structure of scaling limits of large planar maps. *Invent. Math.* **169** 621–670. [MR2336042](https://arxiv.org/abs/MR2336042)
- [38] LE GALL, J.-F. (2010). Geodesics in large planar maps and in the Brownian map. *Acta Math.* **205** 287–360. [MR2746349](https://arxiv.org/abs/MR2746349)
- [39] LE GALL, J.-F. (2013). Uniqueness and universality of the Brownian map. *Ann. Probab.* **41** 2880–2960. [MR3112934](https://arxiv.org/abs/MR3112934)
- [40] LE GALL, J.-F. and MIERMONT, G. (2012). Scaling limits of random trees and planar maps. In *Probability and Statistical Physics in Two and More Dimensions*. *Clay Math. Proc.* **15** 155–211. Amer. Math. Soc., Providence, RI. [MR3025391](https://arxiv.org/abs/MR3025391)
- [41] LE GALL, J.-F. and PAULIN, F. (2008). Scaling limits of bipartite planar maps are homeomorphic to the 2-sphere. *Geom. Funct. Anal.* **18** 893–918. [MR2438999](https://arxiv.org/abs/MR2438999)
- [42] MAO, C., SUN, X. and ZHOU, T. (2014). Inventory accumulation with k products. In preparation.
- [43] MIERMONT, G. (2008). On the sphericity of scaling limits of random planar quadrangulations. *Electron. Commun. Probab.* **13** 248–257. [MR2399286](https://arxiv.org/abs/MR2399286)
- [44] MIERMONT, G. (2013). The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Math.* **210** 319–401. [MR3070569](https://arxiv.org/abs/MR3070569)
- [45] MILLER, J. and SHEFFIELD, S. (2013). Imaginary geometry IV: Interior rays, whole-plane reversibility, and space-filling trees. Preprint. Available at [arXiv:1302.4738](https://arxiv.org/abs/1302.4738).
- [46] MILLER, J. and SHEFFIELD, S. (2015). Liouville quantum gravity spheres as matings of finite-diameter trees. Preprint. Available at [arXiv:1506.03804](https://arxiv.org/abs/1506.03804).
- [47] MULLIN, R. C. (1967). On the enumeration of tree-rooted maps. *Canad. J. Math.* **19** 174–183. [MR0205882](https://arxiv.org/abs/MR0205882)
- [48] POLYAKOV, A. M. (1981). Quantum geometry of bosonic strings. *Phys. Lett. B* **103** 207–210. [MR0623209](https://arxiv.org/abs/MR0623209)
- [49] POLYAKOV, A. M. (1981). Quantum geometry of fermionic strings. *Phys. Lett. B* **103** 211–213. [MR0623210](https://arxiv.org/abs/MR0623210)
- [50] POLYAKOV, A. M. (2008). From Quarks to Strings. Preprint. Available at [arXiv:0812.0183](https://arxiv.org/abs/0812.0183).
- [51] SCHAEFFER, G. (1997). Bijective census and random generation of Eulerian planar maps with prescribed vertex degrees. *Electron. J. Combin.* **4** Research Paper 20, 14 pp. (electronic). [MR1465581](https://arxiv.org/abs/MR1465581)
- [52] SHEFFIELD, S. (2005). Conformal weldings of random surfaces: SLE and the quantum gravity zipper. Preprint. Available at [arXiv:1012.4797](https://arxiv.org/abs/1012.4797).

- [53] SHEFFIELD, S. (2009). Exploration trees and conformal loop ensembles. *Duke Math. J.* **147** 79–129. [MR2494457](#)
- [54] SHEFFIELD, S. and WERNER, W. (2012). Conformal loop ensembles: The Markovian characterization and the loop-soup construction. *Ann. of Math. (2)* **176** 1827–1917. [MR2979861](#)
- [55] SUN, X. and WILSON, D. B. (2015). Sandpiles and unicycles on random planar maps. Preprint. Available at [arXiv:1506.08881](#).
- [56] THORLEIFSSON, G. (1993). Coupling of multiple q -state Potts models to 2d gravity. In *Proceedings of the International Symposium on Nuclear Physics B—Proceedings Supplements* **30** 787–790.
- [57] TUTTE, W. T. (1968). On the enumeration of planar maps. *Bull. Amer. Math. Soc.* **74** 64–74. [MR0218276](#)

DEPARTMENT OF MATHEMATICS
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
77 MASSACHUSETTS AVENUE
CAMBRIDGE, MASSACHUSETTS 02139
USA
E-MAIL: sheffield@math.mit.edu